The Astrophysical Journal, 974:172 (49pp), 2024 October 20 https://doi.org/10.3847/1538-4357/ad6869

Anomaly Detection and Approximate Similarity Searches of Transients in Real-time Data Streams

P. D. Aleo[1,2], A. W. Engel[3], G. Narayan[1,2], C. R. Angus[4,5], K. Malanchev[1,6], K. Auchettl[7,8], V. F. Baldassare[9], A. Berres[1], T. J. L. de Boer[10], B. M. Boyd[11], K. C. Chambers[10], K. W. Davis[7], N. Esquivel[12], D. Farias[4], R. J. Foley[7], A. Gagliano[13,14], C. Gall[4], H. Gao[10], S. Gomez[15], M. Grayling[11], D. O. Jones[16], C.-C. Lin[10], E. A. Magnier[10], K. S. Mandel[11], T. Matheson[12], S. I. Raimundo[4,17], V. G. Shah[1], M. D. Soraisam[18], K. M. de Soto[14], S. Vicencio[12], V. A. Villar[14], and R. J. Wainscoat[10]

[1] Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 W. Green Street, IL 61801, USA; paleo2@illinois.edu
[2] Center for AstroPhysical Surveys, National Center for Supercomputing Applications, Urbana, IL 61801, USA
[3] Pacific Northwest National Laboratory, 902 Battelle Boulevard, Richland, WA 99354, USA
[4] DARK, Niels Bohr Institute, University of Copenhagen, Jagtvej 128, 2200 Copenhagen, Denmark
[5] Astrophysics Research Centre, School of Mathematics and Physics, Queen's University Belfast, Belfast BT7 1NN, UK
[6] McWilliams Center for Cosmology, Department of Physics, Carnegie Mellon University, Pittsburgh, PA 15213, USA
[7] Department of Astronomy and Astrophysics, University of California, Santa Cruz, CA 95064, USA
[8] School of Physics, The University of Melbourne, VIC 3010, Australia
[9] Department of Physics & Astronomy, Washington State University, Pullman, WA 99164, USA
[10] Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, HI 96822, USA
[11] Institute of Astronomy and Kavli Institute for Cosmology, Madingley Road, Cambridge, CB3 0HA, UK
[12] NSF's National Optical-Infrared Astronomy Research Laboratory, 950 North Cherry Avenue, Tucson, AZ 85719, USA
[13] The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, USA
[14] Center for Astrophysics—Harvard & Smithsonian, Cambridge, MA 02138, USA
[15] Space Telescope Science Institute, Baltimore, MD 21218, USA
[16] Institute for Astronomy, University of Hawai'i, 640 N. A'ohoku Place, Hilo, HI 96720, USA
[17] Department of Physics and Astronomy, University of Southampton, Highfield, Southampton SO17 1BJ, UK
[18] NSF's National Optical-Infrared Astronomy Research Laboratory, Gemini North, 670 N. A'ohoku Place, Hilo, HI 96720, USA

Received 2024 March 29; revised 2024 June 19; accepted 2024 July 21; published 2024 October 10

Abstract

We present Lightcurve Anomaly Identification and Similarity Search (`LAISS`), an automated pipeline to detect anomalous astrophysical transients in real-time data streams. We deploy our anomaly detection model on the nightly Zwicky Transient Facility (ZTF) Alert Stream via the ANTARES broker, identifying a manageable ~1–5 candidates per night for expert vetting and coordinating follow-up observations. Our method leverages statistical light-curve and contextual host galaxy features within a random forest classifier, tagging transients of rare classes (*spectroscopic* anomalies), of uncommon host galaxy environments (*contextual* anomalies), and of peculiar or interaction-powered phenomena (*behavioral* anomalies). Moreover, we demonstrate the power of a low-latency (~ms) approximate similarity search method to find transient analogs with similar light-curve evolution and host galaxy environments. We use analogs for data-driven discovery, characterization, (re)classification, and imputation in retrospective and real-time searches. To date, we have identified ~50 previously known and previously missed rare transients from real-time and retrospective searches, including but not limited to superluminous supernovae (SLSNe), tidal disruption events, SNe IIn, SNe IIb, SNe I-CSM, SNe Ia-91bg-like, SNe Ib, SNe Ic, SNe Ic-BL, and M31 novae. Lastly, we report the discovery of 325 total transients, all observed between 2018 and 2021 and absent from public catalogs (~1% of all ZTF Astronomical Transient reports to the Transient Name Server through 2021). These methods enable a systematic approach to finding the "needle in the haystack" in large-volume data streams. Because of its integration with the ANTARES broker, `LAISS` is built to detect exciting transients in Rubin data.

Unified Astronomy Thesaurus concepts: Supernovae (1668); Transient detection (1957); Astronomical methods (1043); Time domain astronomy (2109); Time series analysis (1916); Astrostatistics techniques (1886); Classification (1907); Light curves (918); Random Forests (1935)

1. Introduction

Serendindipity has played a disproportionately large role in breakthrough time-domain astronomical discoveries. For example, the discovery of SN 1987A (Kunkel et al. 1987) in the nearby Large Magellanic Cloud ushered in a renewed interest in supernova (SN) science—spurring systematic searches in the forms of pencil-beam surveys to discover high-redshift SNe for cosmological study (Perlmutter et al. 1997; Schmidt et al. 1998) and rolling searches to drastically increase observed events (Barris et al. 2004). The discovery of SN 1998bw, which is the first SN associated with a gamma-ray burst, was found by astronomers who were observing gamma-ray bursts rather than SNe (Galama et al. 1998). Nearby Type Ia SN 2014J in M82 (Cigar Galaxy) at 3.7 Mpc was discovered by chance during an undergraduate telescope training session at the University of London Observatory by Steve Fossey and his students (Fossey et al. 2014; McIntosh et al. 2014). Recently, Itagaki (2023) discovered SN 2023ixf in M101, the closest SN to Earth since 2014J, from an untargeted search before automated algorithmic discovery.

Beyond those closest and brightest to Earth, transient discoveries of extremely rare origin often arise from astronomers aiming to increase the observable parameter space or find dissimilar objects in ever-increasing data sets (Li et al. 2022). These tasks are worthwhile, because in-depth study of such rare events reveals insights of their progenitors, explosion mechanisms, and diversity—all of which continue to be active areas of research (see, e.g., Margutti et al. 2019; Gagliano et al. 2022; W. V. Jacobson-Galán et al. 2022; Perley et al. 2022; Kuncarayakti et al. 2023; Pierel et al. 2023).

Astronomers have attempted to automate serendipity in a "systematic" manner (see, e.g., Giles & Walkowicz 2019) to increase rare transient discovery. This task is analogous to finding "the needle in the haystack" (see, e.g., Villar et al. 2019), which becomes not only increasingly difficult in the era of large, data-driven surveys (where human vetting for the vast majority of objects is infeasible), but increasingly important due to the limited spectroscopic resources and multiwavelength follow-up observations. It is estimated ∼1% of the approximate 10 million Vera C. Rubin Observatory (Ivezić et al. 2019) Legacy Survey of Space and Time (LSST) transients will be selected for spectroscopic follow-up (Hambleton et al. 2023), down from the ∼10% of all optical transients classified spectroscopically today. Thus, automated algorithms that identify rare or interesting transients must do so with high purity as to be judicious with our resources. Moreover, because it is all but guaranteed that LSST will discover entirely new phenomena, theorized or not (Hambleton et al. 2023; Li et al. 2022), we need adept algorithms that are able to identify events with currently unknown feature distributions.

The astronomical community has been hard at work developing anomaly detection methodologies. Some have focused on static data sets, such as images (Reyes & Estévez 2020; Storey-Fisher et al. 2021; Etsebeth et al. 2024) or spectra (Böhm et al. 2023; Liang et al. 2023), or time-series light-curve data (Rebbapragada et al. 2009; Nun et al. 2014, 2016; Solarz et al. 2017; Giles & Walkowicz 2019; Pruzhinskaya et al. 2019; Soraisam et al. 2020; Webb et al. 2020; Ishida et al. 2021; Lochner & Bassett 2021; Malanchev et al. 2021; Martínez-Galarza et al. 2021; Villar et al. 2021; Perez-Carrasco et al. 2023; Cui et al. 2024). The added complexity in detecting anomalies in time-series data, combined with strong interest in SN photometric classification (e.g., see Lochner et al. 2016; Muthukrishna et al. 2019; Möller & de Boissière 2020; Villar et al. 2020; Qu et al. 2021; Aleo et al. 2023; Gagliano et al. 2023 and references therein), is a key reason why most efforts have concentrated on identifying anomalous light curves.

To that end, most of the effort has been dedicated to full-phase anomaly detection, after the entire light curve has been observed. Only recently has the field focused on real-time anomaly detection (Soraisam et al. 2020; Villar et al. 2021; Muthukrishna et al. 2022; Perez-Carrasco et al. 2023; Gupta et al. 2024), with current efforts combining the SN light curve with contextual information (e.g., host galaxy, spectra, etc.), as seen in Perez-Carrasco et al. (2023) and this work.

In conjunction, instead of discovering transients and isolating anomalies independently, recent effort has been devoted to finding *analogs* of a given object or transient event. Giles & Walkowicz (2019) utilized a similarity score ascribed to the t-distributed stochastic neighborhood embedding (t-SNE; van der Maaten & Hinton 2008) representation of Kepler light-curve features to search for outliers in cluster distributions. Their follow-up work (Giles & Walkowicz 2020) demonstrated the effectiveness of k-nearest neighbor distance in feature-space to efficiently identify anomalous light curves. Martínez-Galarza et al. (2021) expanded on this work by running an unsupervised random forest (Shi & Horvath 2006) on the joint space of Kepler light curves and power spectra with two manifold-learning algorithms (t-SNE and UMAP; McInnes et al. 2018) to create a low-dimensional embedding. They tested their analog-finding ability by analyzing the anomalies' location and clustering in the embedded feature-space, identifying those with shared astrophysical properties. In a different approach, Aleo et al. (2022) simulated bright Zwicky Transient Facility (ZTF) SNe, extracted statistical light-curve features, applied a brute-force k-dimensional tree algorithm to identify nearest neighbors, and visually inspected the closest matches. They discovered 11 previously unreported transient events in ZTF (Bellm et al. 2019) fourth data release (DR4) out of 105 manually vetted objects. Beyond light curves, a similarity search has been used to identify similar galaxy images from their low-dimensional representations (Stein et al. 2021) and associating galaxy images with their optical spectra via cross-modal contrastive learning (Lanusse et al. 2024), as well as selecting synthetic galaxies matching statistical properties of observational data for simulations (Lokken et al. 2023).

Here, we present an approach for simultaneous anomaly detection (real-time and retroactive) and similarity search that uses both SN light-curve and contextual information (host galaxy photometry) from the ZTF Alert Stream. We identify transient events that are of a rare spectroscopic class (*spectroscopic* anomalies), exhibit peculiar behavior or interaction-powered phenomena (*behavioral* anomalies), or are found in host galaxy environments uncommon to their type (*contextual* anomalies), and categorize these broadly as anomalies. In tandem, we find transient analogs across via an approximate nearest neighbors (ANNs) similarity search. We demonstrate the effectiveness of these methods by identifying archival rare or unique transients with high purity and low latency, and reclassifying transients after being prompted to investigate their spectra due to their nearest neighbors' class labels. We call our pipeline Lightcurve Anomaly Identification and Similarity Search (`LAISS`). We provide a schematic overview in Figure 1.

Our paper is structured as follows. In Section 2, we describe our methodology for constructing our database of SNe from the ZTF Alert Stream. In Section 3, we discuss our real-time random forest classifier (RFC) anomaly detection model and the results of its application to the ZTF Alert Stream via the Arizona-NOIRLab Temporal Analysis and Response to Events System (ANTARES) broker (Matheson et al. 2021). In Section 4, we provide additional results based on our model's application to the Young Supernova Experiment (YSE) DR1 (Aleo et al. 2023). In Section 5, we demonstrate the power of low-latency approximate similarity search of SNe, including finding missed SN candidates and the reclassification of some SNe. In Section 6, we speculate on additional applications and extensions of this work, such as performing a calculation and preliminary analysis of our tagged transients' host galaxy masses and star formation rates (SFRs). We conclude in Section 7, and we detail all used features and newly reported SN candidates in Appendix C.

The code is publicly available on Github,[19] and the version of this code used in this work is available on Zenodo.[20]

---

[19] https://github.com/patrickaleo/LAISS-local

[20] doi:10.5281/zenodo.11541806

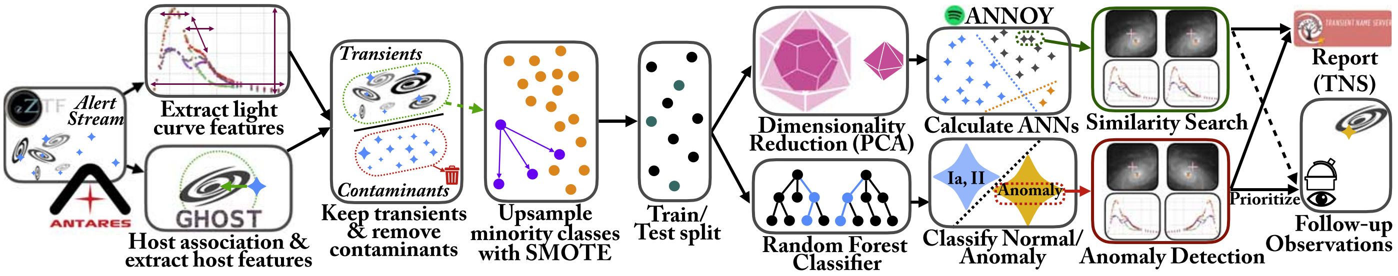

Figure 1. The `LAISS` Pipeline. We ingest ZTF Alert Stream data, and extract statistical light-curve features with `light-curve` (Malanchev 2021) and contextual host galaxy features with `GHOST` (Gagliano et al. 2021). We apply strict preprocessing cuts to remove contaminants and obtain a reference data set of transient candidates. We apply a 70/30 train/test split, subsequently upsampling our training set with SMOTE (Chawla et al. 2002) such that each class is balanced. We train a random forest classifier to perform a binary classification task between "Normal" and an "Anomaly" class. We consider *spectroscopic* anomalies, *contextual* anomalies, and *behavioral* anomalies via expert vetting. In parallel, we apply principal component analysis (PCA) followed by approximate nearest neighbors (ANNs) search in feature-space via `ANNOY` (Bernhardsson 2018; Section 5) to find transient analogs with similar light-curve evolution and host galaxy environment. Finally, we deploy our anomaly detection model onto the ANTARES broker (Matheson et al. 2021) as a *Filter*. For transients with $P(\text{anom}) \geqslant 50\%$ at any point in the light curve, we manually vet them and request for follow-up resources. We report to TNS any new or updated classifications from either our anomaly detection or approximate nearest neighbors search.

2. Constructing Our Training Set

We are motivated by two directives: (1) find anomalous transients in real-time within the ZTF Alert Stream with high purity (see Section 3 for details), and (2) find k ANNs (i.e., analogs) of any transient through our embedding space within ∼ms to scale with the rate of LSST transients. Achieving these directives with current data sets is a strong prototype for eventual application to Rubin transients, where the volume and diversity of transients will demand such infrastructure.

First, we need to construct a reference database of transients (hereafter, "databank"). This section details this process.

2.1. ZTF Data and Broker

ZTF[21] (Bellm et al. 2019) uses the Palomar 48 inch Schmidt telescope equipped with a 47 deg^2 field-of-view camera and an 8 s readout time to observe the entire northern sky in ZTF-g, ZTF-r, and ZTF-i passbands.[22]

Now in Phase II operations, ZTF allocates 50% of camera time and 50% of the spectral energy distribution machine (SEDM) spectrograph time to a 2 night cadence public survey of the entire northern sky. SEDM spectra are uploaded daily to the Transient Name Server (TNS).[23] The Infrared Processing and Analysis Center (IPAC) provides ZTF image reduction and object identification in near real-time, producing transient alerts from raw images in ∼4 minutes. These are available to the community in the public ZTF Alert Stream via alert brokers such as ANTARES[24] (Matheson et al. 2021), Automatic Learning for the Rapid Classification of Events[25] (ALeRCE; Förster et al. 2021), Fink[26] (Möller et al. 2021), Alert Management, Photometry and Evaluation of Lightcurves (`AMPEL`[27]; Nordin et al. 2019), Pitt-Google,[28] and Lasair.[29]

For this work, we extensively use the ANTARES broker. The primary advantage of ANTARES is the allowance of a user-created *Filter*—snippets of Python code, which can analyze and tag light curves with numerical or categorical values. These tags can then be used to filter objects of interest, such as optimal candidates for follow-up observations. For example, our full-phase anomaly detection RFC is implemented as a *Filter* on ANTARES, and has been running since 2023 August 22. More details on this *Filter* are presented in Section 3.1.

ANTARES has processed millions of ZTF objects ("loci"[30]) with their *Filter* `lc_feature_extractor`. This *Filter* uses the `light-curve`[31] package (Malanchev 2021) to quickly (∼ms) calculate 53 statistical light-curve features in each passband. These features and its subset we choose for this work are highlighted in Section 2.3.

Before preprocessing, we start with all ∼2.5 million loci tagged by the `lc_feature_extractor` *Filter*. The overwhelming majority are variable stars, quasistellar objects (QSOs), and active galactic nuclei (AGNs), while a small portion are transient in nature like SNe and tidal disruption events (TDEs). In addition to the desired `light-curve` statistical features, these loci have additional metadata[32] and data products, from R.A./decl. to tags from other science *Filters* and beyond. While no transient loci have associated hosts, we perform our own host associations (Section 2.3.2) using `GHOST` (Gagliano et al. 2021), whose features are used downstream in `LAISS`.

2.2. Preprocessing

We require a highly pure sample of transients to train, test, and validate our models. Although the methods we present here can be extended to variable stars and other phenomena, for this work, we are interested in transients exclusively and thus want to minimize contamination from nontransients.

To obtain a pure transient sample, we impose a strict preprocessing pipeline:

---

[21] http://ztf.caltech.edu

[22] We do not use ZTF-i observations for this work because of poor coverage (∼10% of all ZTF observations) and the 18 month grace period for private survey data before public release, which is not conducive for real-time applications.

[23] https://www.wis-tns.org

[24] http://antares.noirlab.edu

[25] http://alerce.science

[26] https://fink-broker.org

[27] https://github.com/AmpelAstro/Ampel-contrib-sample

[28] https://github.com/mwvgroup/Pitt-Google-Broker

[29] http://lasair.roe.ac.uk/

[30] A locus (plural, loci) is a point on the sky where alerts cluster and is roughly equivalent to an astrophysical object.

[31] https://github.com/light-curve

[32] https://antares.noirlab.edu/properties

1. First, we excise any objects in the galactic plane by requiring $|b| > 15°$. We do so because variable stars far outnumber extragalactic transients along the galactic plane, and the light curves of longer period variables (e.g., Mira stars) can resemble those of SNe in statistical feature-space (Malanchev et al. 2021; Aleo et al. 2022). This selection cut leaves ∼1 million objects.
2. We require both ZTF-g and ZTF-r `light-curve` features, which are only calculated for bands with at least four observations. Thus, we remove all objects with fewer than four observations in either passband, leaving ∼0.5 million objects.
3. At this stage, we have yet to distinguish between variable stars, asteroids, bogus observations, AGNs, transients, and the like. Because ANTARES has no built-in star/galaxy or variable/transient separator, we outsource this step to existing and proven infrastructure.

   We query the ALeRCE light-curve classifier (Sánchez-Sáez et al. 2021) and their stamp classifier (Carrasco-Davis et al. 2021). We prioritize the results from the light-curve classifier, which parses objects into 15 subclasses of variable and transient objects across "Transient," "Stochastic," and "Periodic" sources. We keep only the objects that have a Transient classification from light-curve classifier ("SN Ia," "SN Ibc," "SN II," "SLSN"), or stamp classifier ("SN") as its most probable classification, prioritizing the results from the light-curve classifier.[33] After cuts, we have ∼10,000 objects, the majority of which are bona fide transients.
4. From this sample, we obtain host galaxy features using the `GHOST` software[34] (Gagliano et al. 2021). `GHOST` is a database of 16k PanSTARRS (Pan-STARRS1, hereafter PS1; Chambers et al. 2016) spectroscopic SNe and the catalog-level properties of their host galaxies, equipped with analysis tools for transient-host association. A final association is made using the directional light–radius (DLR) method at the catalog level. We do not opt to use the gradient ascent method at the postage-stamp level[35] for computational reasons.

   Because we require host galaxy features as part of our input data schema, transients with unsuccessful host associations (e.g., hosts not in the PS1 host galaxy catalog, transients occurring in faint host galaxies that are either not present in the PS1 catalog or cannot be found during the association etc.) are discarded. From 9402 bona fide transients, 6571 have successful host associations. However, we additionally remove transients whose host galaxies have null values for any of the features presented in Appendix B. This requires PS1 photometry in *grizy* passbands.[36] After cuts, we have 5837 remaining objects for which we query TNS for a spectroscopic classification (if it exists).

Table 1
Preprocessing Cuts

Selection Cuts	No. Loci Remaining
Tagged by `lc_feature_extractor` Filter	2,536,582
Have $|b| > 15°$	976,842
Require both ZTF-g, ZTF-r passbands	488,300
Have ALeRCE "SN" classifications	9420
Have successful `GHOST` host galaxy association	6571
Have nonnull host galaxy features	5837
Cross-matched to resolved extended source (PS1-PSC)	5472

Note. The number of loci (representing potential transient candidates) after selection cuts to construct our databank.

5. As a final attempt to increase the purity of bona fide transients in our databank, we query the PS1 point-source catalog (PS1-PSC; Tachibana & Miller 2018). The PS1-PSC contains ∼1.5 billion sources from the PS1 DR1, and is used within ZTF's real-time extragalactic alert stream to automatically reject stellar sources.

   We query the PS1-PSC at two locations per transient candidate: the location of the transient, and the location of the host galaxy, both with a search radius of 1″. We retain transient candidates whose host galaxy locations match to a resolved extended object with >50% probability, and whose transient locations either (A) also match to resolved extended object with >50% probability (in the case of a transient occurring at/near the galaxy center) or (B) match to no existing counterpart in the PS1-PSC catalog (in the case of a transient occurring far/offset from the galaxy center). After cuts, we have a remaining data set of 5472 likely bona fide transients.

A table summarizing the counts of loci at each stage of our preprocessing and quality cuts is found in Table 1.

Of our remaining 5472 objects, only 1656 (≈30%) have a spectroscopic classification available from TNS. This leaves 3816 transients with no spectroscopic classification. The spectroscopic breakdown is found in Table 2. The subsequent process for generating our ZTF training set from this full databank is explained in Section 2.4.

2.3. Feature Selection

Here, we outline our motivation for feature selection. Instead of using the observed, irregularly sampled light curve, we opt to use derived statistical features for the anomaly detection and similarity search methods (Sections 3.1, 5.1). A review of the challenges astronomers face when analyzing light curves of astronomical sources and the practices for their characterization via statistical feature extraction is found in Babu & Mahabal (2016).

The main advantage of using statistical features is that they transform the sparse, gappy, and heteroscedastic light curves into regularly sampled representations, enabling the use of standard time-series methods. These methods do not rely on imputation or interpolation; they are calculated only from the data that is observed. The fast, efficient, and scalable calculation of statistical light-curve features is now ubiquitous in LSST brokers. The `light-curve` package (Malanchev 2021), which is used to extract the light-curve features for our data set, is already integrated into the pipelines of ANTARES

---

[33] According to Carrasco-Davis et al. (2021), the stamp classifier has an overall 87% SN completeness of SN, and agrees with 78% of the light-curve classifier SN.

[34] https://pypi.org/project/astro-ghost/

[35] https://ps1images.stsci.edu

[36] This effectively removes the faintest sources (such as those above the signal-to-noise ratio limit in only a few bands) and extremely red ones detected only in the reddest filter (y). This includes the faintest host galaxies (in which SLSN are preferentially found) and red sources like brown dwarfs and high-redshift quasars (Magnier et al. 2020). The impact on this selection function is left unquantified and left for future work.

Table 2
Databank Spectroscopic Breakdown (5754 Objects)

Misc. (3852)	SN Ia (1160)	SN II (340)	SN Ib/c (75)	Exotic (45)
Phot (3816)	Ia (1098)	II (243)	Ib (25)	TDE (20)
CV (8)	Ia-91T-like (37)	IIn (59)	Ic (21)	SLSN-II (14)
AGN (7)	Ia-pec (10)	IIb (21)	Ic-BL (14)	SLSN-I (11)
SN I (6)	Ia-91bg-like (6)	IIP (14)	Ibn (7)	…
SN (5)	Ia-CSM (4)	II-pec (2)	Ib/c (5)	…
Other (5)	SN Iax[02cx-like](3)	IIn-pec (1)	Ib-pec (2)	…
Nova (3)	Ia-SC (2)	…	Icn (1)	…
Varstar (2)	…	…	…	…

Note. The spectroscopic classifications from TNS used in our databank before a train/test split and upsampling.

(Matheson et al. 2021), AMPEL (Nordin et al. 2019), and Fink (Möller et al. 2021) broker systems.[37] Moreover, it has been selected for use with LSST. The Rust version of `light-curve` package outperforms other implementations by a factor 1.5–50, such that it only takes a few milliseconds per CPU core to extract "cheap" features for 1000 light-curve observations (Lavrukhina & Malanchev 2023).

Although relying solely on interpolation-based methods like Gaussian process (GP; Rasmussen & Williams 2005) is a popular choice (see, e.g., Lochner et al. 2016; Boone 2019; Alves et al. 2022), these methods tend to be slower and are not ubiquitous in broker pipelines. Generally, light-curve interpolation, smoothing, and padding are sometimes necessary or at least tend to improve results in a variety of deep-learning applications: multilayer perceptrons (Demianenko et al. 2023), normalizing flows (Demianenko et al. 2023), temporal convolutional networks (Muthukrishna et al. 2019), recurrent neural networks (Charnock & Moss 2017; Möller et al. 2021; Gagliano et al. 2023), convolutional neural networks (Pasquet et al. 2019; Qu et al. 2021; Burhanudin & Maund 2023), Bayesian neural networks (Demianenko et al. 2023), variational autoencoders (VAEs; Villar et al. 2020; Boone 2021; Aleo et al. 2023), and transformers (Donoso-Oliva et al. 2023; Moreno-Cartagena et al. 2023; Pimentel et al. 2023). Oftentimes, the physical interpretation of such learned features from latent or deep-learning representations can be nebulous.

In the literature, there has been proven successes of using manually selected statistical features for classification (e.g., the ALeRCE light-curve classifier, Sánchez-Sáez et al. 2021; *dmdt*-mappings, Mahabal et al. 2017; anomaly detection $p(dm|dt)$ distributions, Soraisam et al. 2020; the `SNAD Transient Miner`, Aleo et al. 2022). The lessons learned from these works inform our feature selection, methods, and approach to follow.

2.3.1. Statistical Light-curve Features

We start with the suite of 106 light-curve features (53 in ZTF-g, 53 in ZTF-r bands) automatically extracted (at the last observed epoch; features calculated at earlier phases are overwritten) from the `lc_feature_extractor` *Filter* in ANTARES using the `light-curve` package (Malanchev 2021). This ensemble of features describe different aspects of the light-curve shape, capturing rise information, error distribution, periodic signals, skewness, etc. Moreover, transients whose feature values are in the tails of their parent feature distribution are deemed outliers and could indicate anomalous activity.

We know from Malanchev et al. (2021; their Figure A3) that many of the calculated light-curve features are correlated. To mitigate the effect of correlated features and attempt to partially reduce the dimensionality of our data set, we drop features that are *not* used in Aleo et al. (2022), leaving us with 82 light-curve features (41 per band). Because Aleo et al. (2022) were able to successfully find previously undiscovered SNe candidates among nearly ∼1 million sources (most of which were standard variables) using this feature set, we were confident that if we adopted a similar feature set we would also be successful given that the data set we work with here has significantly higher purity of transients. Lastly, we removed 10 features (per band) that were strongly dependent on the cadence (or time-dependence) of observations. We do this because downstream we want to find analogs of the objects themselves without biasing *how* their light curves are measured. This includes features pertaining to measuring the period or calculating the periodogram.

Our final light-curve feature list comprises 62 total: 48 magnitude-based and 14 flux-based features, all of which are brightness related; see the complete list of features and their description in Appendix A.

2.3.2. Host Galaxy Features

As SN discovery rates grow at a rapid pace, particularly in today's expansive landscape of systematic sky surveys, the correlation between SNe and their global galactic environments has started to crystallize (e.g., Li et al. 2011; Qin et al. 2022, and references therein). For instance, the rate of core-collapse (CC) SNe, whose progenitors are short-lived (<50 Myr) massive stars, is tightly coupled to the SFR of a galaxy (Graur et al. 2017). Similarly, there are unique galactic correlations that can be gleaned for many SN demographics, and this information should be considered when quantifying the "anomalous" nature of a transient. Consider the following: all but a single SN Iax to date have been discovered in late-type galaxies (Lyman et al. 2018). Thus, a new discovery of a second SN Iax in an early-type galaxy could be considered *more* anomalous than the same SN Iax (i.e., with the same light-curve evolution and derived properties) in a late-type galaxy.

Our work, which incorporates derived global host galaxy properties, lays the groundwork for such future analyses of varying SNe types. Thus, the object's anomaly score is affected by not only its light-curve behavior, but its global host galaxy properties as well. We use approximately a 1:1 ratio of light-curve features to host galaxy features.[38] Ultimately, global host galactic environments of the SNe are considered to be meaningful, although weighted less when compared to unusual light-curve behavior. As for local correlations, we conduct no formal investigation because of the general lack of spatially resolved explosion site data.

---

[37] As for non-`light-curve` implementations, ALeRCE has their own in-house feature extractor (Förster et al. 2021).

[38] In practice, this weighting does not hold exactly (e.g., see Figure 6, where 8 of the top 10 most important features used in our anomaly detection task are light-curve features, and the remaining two are host galaxy features).

Incorporating contextual information into SN identification and classification tasks has steadily increased in the last decade. Foley & Mandel (2013) demonstrated that host galaxy morphology and color were sufficient to construct photometric SN Ia samples at the purity level of then-current light-curve methods. Baldeschi et al. (2020) used photometric host features (colors and moments) from the PS1 DR2 catalog to build an RFC and assign probabilities between low and highly star-forming host galaxies, subsequently to help distinguish SNe II from SNe Ia. ALeRCE developed a postage-stamp classifier (Carrasco-Davis et al. 2021)[39] to disambiguate SNe from non-SNe (AGN, variable stars, asteroids, bogus detections) with a single detection image, in part relying on the fact that many SN detections contain a visible host galaxy in both science and reference images, but only the science and difference images contain the flux from the SN. Gagliano et al. (2021), which debuted `GHOST`, found that photometric features of transient-hosting galaxies alone were sufficient in classifying SNe II and SNe Ia with ∼70% accuracy when fed into an RFC. Gagliano et al. (2023) extended such work to combine light-curve and host galaxy information for adaptive, real-time photometric classification, with emphasis on early time phases, achieving an accuracy of 82% ± 2% within 3 days of an event's discovery.

Moreover, publicly accessible SN data products have been keen to include host galaxy information: the YSE DR1 (Aleo et al. 2023) included the associated PS1 host galaxies of 1975 SNe, and the simulated Extended LSST Astronomical Time-Series Classification Challenge (ELAsTiCC; Narayan & ELAsTiCC Team 2023) produced LSST alert packets with both transient photometry and host information. In the LSST era, where transient spectroscopic classification capabilities will be capped at approximately 1% of all transients (Hambleton et al. 2023), any additional information to augment the sparse transient photometry will be required to fully characterize the SNe and its environment. It is sensible for host galaxy information to provide such a link.

With the laundry list of works that were successful in SN identification and classification tasks using host galaxy information, as well as its prevalence in data products, it is a natural extension to use it in the context of anomaly detection and similarity search.

`GHOST` (Gagliano et al. 2021) natively connects our transients to a host in the PS1 galaxy catalog and its hundreds of catalog-level properties.[40] All host galaxy features per transient are retrieved from the PS1 catalog or explicitly calculated with one host association calculation, and are reused downstream in our anomaly detection and similarity search methods (we make the reasonable assumption that the host association and derived host properties of each SNe are constant in time).

The selection of global host features includes the Kron (1980) flux from g, r, i, z, y filter stack detections, the normalized offset of the SNe from its host with respect to the DLR, and the extendedness measure for the g, r, i, z, y filter stack detections based on the deviation between point-spread function (PSF) and Kron (1980) magnitudes normalized by the PSF magnitude uncertainty, to name a few. Our final host galaxy feature list comprises 58 total. See the full list and a brief description of each host galaxy feature used in this work in Appendix B, and discussion of our feature correlations in Appendix C.

2.4. ZTF Training Set

We build our training set using our databank presented in Table 2. Naturally, due to preexisting follow-up strategies and relative SN rates, our spectroscopic data set is heavily class-imbalanced; our majority classes are vastly overrepresented, and our minority classes are vastly underrepresented. Moreover, any algorithm to demarcate anomalies from normal SNe trained on such a class-imbalanced data set would incur a bias toward the majority classes.

Because we want an algorithm to learn the full feature-space of these different SN classes, as well as to mitigate any bias incurred from class-imbalance, we will upsample all classes. However, before upsampling, we restrict our SNe classes to those that have at least 14 members (before train/test split). This strikes a balance between having several rare SNe classes in our training set while retaining enough members in each class to properly upsample in feature-space (via interpolation between known examples) after a train/test split. In this way, the feature-space can be satisfactorily representative of the full, unknown feature distribution to achieve sufficient performance.

There are 3816 photometric objects without a spectroscopic classification ("Phot" class). Those that do have spectroscopic classifications breakdown into 1098 SN Ia, 243 SN II, 59 SN IIn, 37 SN Ia-91T-like, 25 SN Ib, 21 SN IIb, 21 SN Ic, 20 TDE, 14 SN Ic-BL, 14 SN IIP, and 14 SLSN-II. We perform a standard 70/30 train/test split. The training set before upsampling breaks down into 2656 photometric-only ("Phot" class), 777 SN Ia, 175 SN II, 44 SN IIn, 28 SN Ia-91T-like, 19 SN Ib, 15 SN IIb, 12 SN Ic, 12 TDE, 9 SN Ic-BL, 9 SN IIP, and 11 SLSN-II. From this class-imbalanced training set, we upsample to a class-balanced training set using the synthetic minority oversampling technique (SMOTE; Chawla et al. 2002), such that each minority class has the same number of samples as the majority class (2656), totaling 31,872 objects. Thousands of labeled samples per class well encapsulate the feature-space needed by the RFC for training.

SMOTE works by drawing random samples along vectors joining every grouping of k objects in feature-space until all classes are balanced. We use $k = 8$ neighbors; this is the largest value we can use for our training set, because our smallest class to be upsampled for training has nine members (SN IIP, SN Ic-BL). Note that these upsampled features are derived from the full light curve (i.e., not partial light curves) and the host galaxy features. Our final training set is detailed in Section 3.1.

We leave the spectroscopic test set as is, and do not perform upsampling because we want to evaluate our performance on a representative sample of true events, matching the selection function we impose. Our spectroscopic test set is composed of 321 SN Ia, 68 SN II, 15 SN IIn, 9 SN Ia-91T-like, 6 SN Ib, 6 SN IIb, 9 SN Ic, 8 TDE, 5 SN Ic-BL, 5 SN IIP, and 3 SLSN-II. Our observed photometric set is not used in training or testing, but is used in our ANNs search (Section 5).

Note that, at this stage, we do not perform any linear or nonlinear transformation of feature values. We do no rescaling for our RFC (Section 3.1), but we do standardize features by removing the mean and scaling to unit variance for principal component analysis (PCA) before our ANN search with `ANNOY` (Section 5). Due to the limited size of our

---

[39] Note that the postage stamp was centered on the detection, thereby encoding *some* local/offset information.

[40] https://outerspace.stsci.edu/display/PANSTARRS/PS1+Database+object+and+detection+tables

spectroscopically labeled data set, especially for rare classes with few members, we opt for a train/test set instead of the more commonly used training/test/validation set.

SMOTE (Chawla et al. 2002) is known to have certain limitations, particularly when dealing with small training sets and high-dimensional feature-spaces. It is prone to generating synthetic points that lie along the lines connecting original data points rather than curves. This behavior, while expected, indicates that SMOTE may not fully capture the complex, curved manifolds (Bellinger et al. 2016) inherent in our 120-dimensional feature-space and the complex phenomena of SNe light curves and hosts.

Given these constraints, we acknowledge that alternative oversampling techniques, such as the adaptive synthetic sampling method (He et al. 2008) and normalizing flows (Papamakarios et al. 2021), might better address these issues by considering the local density and distribution of data points, or correlations across features. This is a subject of future work.

Despite these limitations, our use of SMOTE facilitates direct comparisons with other anomaly detection methods, such as the FLEET algorithm (Gomez et al. 2020, 2023a, 2023d), which also employs SMOTE in a similar context.

3. Anomaly Detection in ZTF Alert Stream

With any anomaly detection task, the definition of "anomaly" and its associated evaluation metric must be clear. For this work, we consider a transient to be an anomaly if it falls into any of the following three categories:

1. *Spectroscopic anomaly*. This is any transient whose spectroscopic label is *not* of Type Ia-normal, Ia-91T-like,[41] II-normal, and IIP (the "Normal" type of our RFC). Thus, we consider the remaining transient labels explicitly defined in our training set (e.g., TDE, SLSN-II, SN IIn, SN IIb, SN Ib, SN Ic, SN Ic-BL) as well as labels outside our training set (e.g., SN Ia-91bg-like, SN Iax, SN Ibn, SN Icn, SLSN-I, etc.) as anomalous (the "Anomaly" type of our RFC). Note that, when we consider spectroscopic anomalies alone, we do so using the original TNS classification label without looking at the spectra. However, when we consider spectroscopic anomalies after vetting (grouped with contextual and behavioral anomalies), we determine as such through the reclassification of the transient's spectra with `SNID` (Blondin & Tonry 2007; if needed). This is often spurred by "Anomaly" classifications from our model on objects originally outside of our anomalous classification taxonomy, or objects whose ANNs are instances of the "Anomaly" class.
2. *Contextual anomaly*. This is any transient whose host environment is peculiar (e.g., an SN II-normal in an evolved elliptical galaxy,[42] or an SN in a rare galaxy type) based on our current understanding of host galaxy correlations. Note that we do not consider SNe Ia in spiral galaxies to be rare enough alone to be considered anomalous. We identify contextual anomalies through manual vetting of the host galaxy type in conjunction with the known or likely SN type.
3. *Behavioral anomaly*. This is any transient whose light-curve behavior exhibits an atypical or peculiar evolution indicative of an underlying physical process (e.g., a prominent second bump in the light curve caused by circumstellar interaction) and not an observational effect (e.g., missing rise information, color information, imaging artifacts, etc.). We identify behavioral anomalies through manual vetting by an expert of the light-curve evolution with respect to the typical evolution of the SN type.

Our philosophy in constructing `LAISS` to discover anomalous transients is similar to that of FLEET (Gomez et al. 2020) to find SLSNe (Gomez et al. 2023a) and TDEs (Gomez et al. 2023d), in part due to their proven success. Our guiding principles are as follows:

1. We "classify" only anomalies with no regard for the classification success of other transients.
2. We favor anomaly sample purity over sample completeness.
3. We prioritize low latency and simple model construction to allow for real-time identification.
4. We aim for a manageable list of currently active flagged anomalous transients, such that an expert can manually vet the source and select the best targets each night.
5. We flag anomalies while they are bright to enable photometric and spectroscopic follow-up.

3.1. Full-phase Anomaly Detection with Random Forest Classifier and Isolation Forest

3.1.1. Spectroscopic Anomalies (Before Vetting)

With our class-balanced training set, we remove objects with no spectroscopic label (the "Phot" class), leaving us with 29,216 labeled spectroscopically confirmed and SMOTE-generated events. From our testing set, we assign SN Ia, SN Ia-91T-like, SN II, and SN IIP as "Normal" (403 objects), and TDE, SN IIn, SLSN-II, SN IIb, SN Ib, SN Ic, and SN Ic-BL as "Anomaly" (52 objects).

Now, we set out to "classify" anomalies. We do not classify by transient class label; rather, we do a binary classification into "Normal" and "Anomaly" classes. To discourage our algorithm from tagging normal SNe Ia as anomalous, we assign SN Ia-normal and SN Ia-91T-like as "Normal," despite the latter being a more rare overluminous subtype. We find that including SN Ia-91T-like as "Anomaly" resulted in a higher rate of Ia tagged as anomalous.[43] This is perhaps due to the large photometric similarity in the light-curve profile between SN Ia-91T-like and SN Ia-normal (unlike SN Ia-91bg-like) because distinguishing factors like absolute magnitudes are not explicitly used. Thus, to be consistent throughout this work, we consider SN Ia-91T-like to be like any other SN Ia-normal, and state that our model is not well suited to tag SN Ia-91T-like as "Anomaly."

We perform many iterations of varying machine learning algorithms for anomaly detection while varying hyperparameters via grid search. Some of these tests included a VAE, support vector machine (SVM), isolation forest (IF), and an

---

[41] Our model performs more competitively when we consider the rare subtype SN Ia-91T-like SNe as nonanomalous due to overlap in feature-space with SN Ia-normal. See Section 3.1 for details.

[42] Irani et al. (2022) concluded approximately 0.3% of all core-collapse (CC) SNe occur in elliptical galaxies, as derived from the spectroscopically complete ZTF-BTS.

[43] We find that including SN Ia-91T-like as "Anomaly" resulted in a higher rate of Ia tagged as anomalous, lowering our purity about 10%.

Table 3
Comparison of Anomaly Detection Algorithms

Algorithm	Hyperparameters via Grid Search	Recall (Anomaly) (%)	Purity (Anomaly) (%)	Total Anomalies Predicted (Test Set)	False Positive Characteristics
VAE	batch_size = {1, 64, 128, **256**, 512} epochs = {5, 25, **50**, 75, 100} intermediate_dim = {2, 8, 16, **64**} latent_dim = {1, **3**, 5, 10} pos_weight = {0.75, **0.9**, 0.95, 1, 2, 5} metrics = {**Precision**(), Recall(), Accuracy()} activation = **“relu,” “sigmoid”** optimizer = **“adam”**	6	10	30/455 (7%)	Bright ($m \lesssim 18$ mag), SN Ia in spiral/irregular hosts SN II with r plateau, linear g decline
SVM	kernel = {**“rbf,”** “linear,” “poly”} C = {0.5, **1**, 5, 10} gamma = {“scale,” **“auto”**} decision_function_shape = {**“ovo,”** “ovr”} class_weight = **“balanced”**	25	19	68/455 (15%)	Often in edge-on, faint, or compact hosts
IF	n_estimators = {50, **100**, 500} max_features = {5, 15, 25, **35**, 45} contamination = **13**	23	20	60/455 (13%)	Very bright ($m \lesssim 16.5$ mag)
RFC	n_estimators = {50, **100**, 500} max_depth = {5, 15, 25, **35**, 45} max_features = {5, 15, 25, **35**, 45} class_weight = **“balanced”**	29	52	29/455 (6%)	Large gaps in light curve, or long-lived

Note. Comparing tested anomaly detection algorithms with grid search. Bold text indicates the hyperparameters chosen for the final algorithms from a grid search. The “Total Anomalies Predicted” column indicates how many objects were tagged as “Anomaly” out of the test set sample of 455 objects for each tested algorithm. The corresponding percentage is shown in parentheses. The “False Positive Characteristics” column indicates common patterns of “Normal” SNe tagged as anomalous, to reveal common failure modes. After comparing several algorithms, the random forest classifier performed the best, with highest overall anomaly purity balanced with a suitable number of anomalies tagged relative to their estimated population.

RFC. A comparative table summarizing the performance metrics of all tested algorithms is found in Table 3. We highlight the combination of hyperparameters from the grid search resulting in the best anomaly purity, while balancing both the recall and fraction of predicted test set anomalies. Additionally, we note commonalities among truly nonanomalous SNe classified as anomalous (“false-positive”). A more in-depth modification of standardized out-of-the-box algorithms (e.g., customizing a weighted reconstruction class loss or anomaly specific class loss using pairwise distances in latent space) for specialized anomaly detection tasks is a subject of future work.

Ultimately, we achieved the best and most robust performance with the highest sample purity on the spectroscopic test set via the `sklearn` implementation of an RFC.[44] An RFC is an ensemble learning method constructed of many decision trees (each trained on bootstrap samples using a random subset of features without replacement), which outputs the mode of the classes in a classification task.

We use an RFC with the following hyperparameters: 100 trees (`n_estimators = 100`), a tree depth of 35 (`max_depth = 35`), a maximum of 35 features out of the 120 input features per tree (`max_features = 35`), a contamination level of 13%[45] (`contamination = 0.13`), and balance the weighting of normal to anomalous classes (```class_weight = ``balanced''```). Surprisingly, increasing the weighting toward anomalous classes had no statistically significant increased performance, and thus, we kept the simpler balanced weighting. We optimized our hyperparameters with a grid search: `n_estimators` from 50 to 500 in steps of 50, `max_depth` from 5 to 65 in steps of 5, `max_features = 35` from 5 to 65 in steps of 5, and `class_weight` from 1 to 10 in steps of 1. We used the Gini impurity as our measure for the quality of feature split. To estimate the classifier’s uncertainties, we run each version of the model 25 times using different random seed initializations. Training for each run was performed on a 2 GH. Quad-Core Intel Core i5 (macOS Version 11.7) and finished in ∼20 s after grid search.

The simplest overview metrics to understand the performance of a classifier are completeness, purity, and accuracy. These metrics are defined for a single class as

$$\begin{aligned} \text{Completeness} &= \frac{\text{TP}}{\text{TP} + \text{FN}} \\ \text{Purity} &= \frac{\text{TP}}{\text{TP} + \text{FP}} \\ \text{Accuracy} &= \frac{\text{TP} + \text{TN}}{\text{S}} \end{aligned} \tag{1}$$

where TP (FP) is the number of true (false) positives, TN (FN) is the number of true (false) negatives, and S is the total sample size.

---

[44] `sklearn.ensemble.RandomForestClassifier`, see https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html.

[45] To match our observed value within a conservative limit to account for potential misclassifications and peculiar behavior.

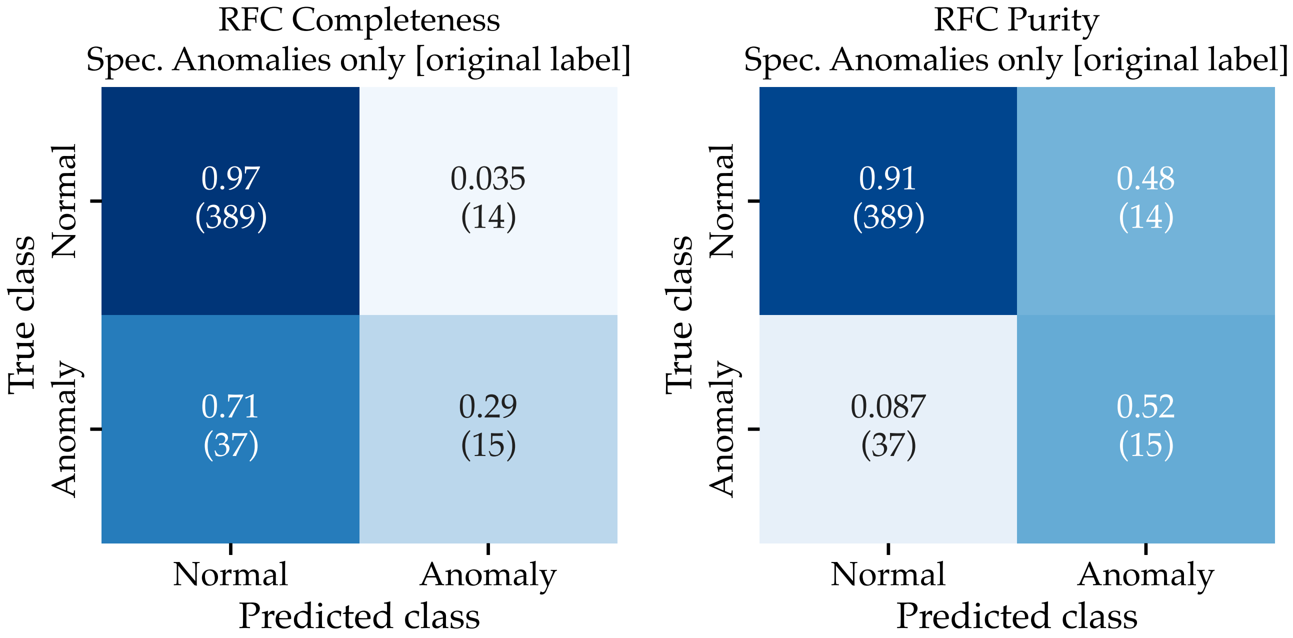

Figure 2. Completeness (left panel) and purity (right panel) confusion matrices of our spectroscopic test set for binary classification performance between "Normal" SNe (SN Ia-normal, SN Ia-91T-like, SN II-normal, SN IIP) and "Anomaly" SNe (TDE, SLSN-II, SN IIn, SN IIb, SN Ib, SN Ic, SN Ic-BL), based on the original TNS spectroscopic label only (i.e., *spectroscopic* anomalies before vetting). We achieve 52% purity of detecting *spectroscopic* anomalies alone.

In this work, completeness ("recall") quantifies the percentage of a true type that is correctly classified ("Normal" or "Anomaly"). Purity ("precision") quantifies the percentage of a predicted type that is correctly assigned the true type. Accuracy is the overall fraction of events that were correctly classified into their respective groups.

The performance of this RFC on the spectroscopic test set is shown as completeness (left) and purity (right) confusion matrices in Figure 2. Note that for these confusion matrices we only consider the value of the original spectroscopic label from TNS alone, and do not yet consider expert vetting for updated classifications, peculiar behavior, or atypical host galaxy environment; such analysis can be found later in this section.

Our RFC achieves an overall accuracy of $89\% \pm 1\%$ with a $97\% \pm 1\%$ "Normal" completeness and $91\% \pm 1\%$ "Normal" purity, with $29\% \pm 3\%$ "Anomaly" completeness and $52\% \pm 5\%$ "Anomaly" purity (with 1σ uncertainty derived from the different random seed iterations).

In a magnitude-limited survey like ZTF, we can expect about 70% of observed SNe to be normal SNe Ia, and about 20% to be normal SNe II, leaving about 10% to be of a rare/anomalous nature (for ZTF-BTS, the measured relative rates are ~72% SNe Ia, ~20.5% SNe II, ~5.5% SNe Ib/c, ~2.5% SLSNe; see Fremling et al. 2020). Indeed, in our spectroscopic test sample of 455 SNe, 52 (11%) are considered anomalous. To have a completeness of 29% and a purity of 52% on the remaining 11% of our test set implies that our RFC model significantly outperforms random selection. Because our test set is heavily imbalanced, the probability of randomly selecting "Anomaly" should be proportional to its prevalence in the data set. Thus, if we expect the observed sample to be roughly 10% "Anomaly," we would expect to successfully recover roughly 10% anomalous events by random selection. Moreover, at 29% completeness, our RFC still manages to identify roughly 1/3 true anomalies, a significant recovery fraction.

We also find that our RFC model vastly outperforms a more sophisticated baseline model commonly used for anomaly detection (see Table 3): an IF (Liu et al. 2012). An IF is an unsupervised decision-tree based method that isolates outliers by randomly selecting a feature from an array of features, and randomly selecting a threshold value (between the max and min) to split. This random partitioning of features and values produces shorter paths in trees for outlier/anomalous data points when aggregated in a tree ensemble. We use the same hyperparameter setup when available (`n_estimators = 100`, `max_features = 35`, `random_state = 11`) and a contamination level of 13% (`contamination = 0.13`). If we compare the 29 objects tagged "Anomaly" (where 15 are correctly recovered, according to the spectroscopic label only) to the 29 objects with the highest anomaly score from the IF, we find the IF only successfully recovers seven events. Moreover, 27/29 objects deemed most anomalous from the IF are brighter than 18.5 mag, with 13/29 exceeding ~16.5 mag in brightness. Of those, 11/13 are simply nearby, normal SNe Ia, and only 2/11 are what we consider anomalous, SNe Ic. In essence, the most anomalous objects identified by the IF are simply bright, nearby, normal SNe that can be found with a selection cut. However, this result may be unsurprising given the disparity between a supervised approach like our RFC model compared to an unsupervised IF model.

A table highlighting all 29 test set objects tagged anomalous by the RFC model can be found in Table 4. First, we consider the resulting distributions from the spectroscopic label alone before expert vetting and potential reclassifications. The objects of "Normal" spectroscopic classes are nine SN II (13.2% of all SN II in test set) and five SN Ia (1.5% of all SN Ia in test set). Similarly, for our anomalous classes, our model tagged six TDE (75%), three SN IIn (20.0%), two SLSN-II (66.7%), one SN IIb (16.7%), one SN Ib (16.7%), one SN Ic (11.1%), and one SN Ic-BL (20.0%). However, because of small number

Table 4

The 29 Events from Our Spectroscopic Test Set with P(anom) $\geqslant 50\%$ at *Full* Light-curve Phase, Ordered by P(anom)

IAU Name	TNS Class	New Class?	z	Peak $M_{\rm abs}$	P(anom)	Remarks
2020yue	SLSN-II	TDE[a]	0.204	−21.8	91	Visible for ∼220 days. Blue, nuclear.
2021uzt	SN II[b]	SN IIn	0.061	−19.4	89	No visible host (host association incorrect).
2018iih	TDE	⋯	0.212	−21.3	76	Visible for ∼1000 days. No data at peak.
2021iui	SN IIn	⋯	0.106	−19.7	76	Faint host.
2020aeuh	SN Ia	SN Ia-CSM[c]	0.126	−19.8	72	Second bump (likely CSM interaction). Visible for ∼175 days.
2018khn	SN IIn	⋯	0.091	−19.4	71	Faint host. Visible for ∼450 days.
2021bxq	SN IIn	⋯	0.09	−19.0	70	Faint host. Visible for ∼225 days.
2021gje	TDE	⋯	0.358	−21.9	69	Visible for ∼50 days. Blue, nuclear.
2021aazw	SN Ic-BL	⋯	0.024	−16.6	68	Visible for ∼50 days.
2021our	SN Ia	⋯	0.039	−18.2	66	Normal SN Ia. Tagged because underluminous?
2021aff	SN II	⋯	0.051	−17.3	64	Incorrect host association. No decline. Visible for ∼150 days.
2021fyp	SN II	⋯	0.053	−18.6	64	Peculiar light curve. Spectrum has He, weak H (possible IIb).
2021cpi	SN II	SN IIb	0.055	−18.2	64	Lack of strong features in first peak spectrum (shock cooling?).
2021ackd	SN Ia	SN Ia-91bg-like	0.050	−18.1	63	Matches to SN 1991bg, 2007ax, 1986G.
2021ckb	SN II	SLSN (He-rich)	0.070	−19.1	62	Narrow H_α emission from host. Match to PTF10hgi.[d]
2021nwa	TDE	⋯	0.047	−18.2	61	Blue, nuclear. Visible for ∼150 days.
2020tnq	SN Ia	⋯	0.033	< -18.3	61	First observations around +26 days after peak (from spectrum).
2021M	SN IIb	⋯	0.011	−15.2	60	Edge-on host galaxy. Visible for ∼55 days.
2021mhg	TDE	⋯	0.073	< -19.5	60	Rebrightening/second transient event at location.[e]
2020abah	SN II	⋯	0.03	−16.5	60	Known member of long-rising SN II class at ∼90 days.[f]
2021adpx	SN Ia	SN Ia-91bg-like	0.047	< -18.7	59	∼20 day r-band gap. Spiral host. Matches to SN 2002cf, SN 2006gt.
2020ywo	SN II	SN Ia-CSM	0.0475	< -18.8	58	Asymmetric H_α hints IIn, but redwards fit match Ia.
2020rmk	SN II	⋯	0.025	−16.9	57	Candidate member of long-rising SN II class.
2020acka	TDE[g]	⋯	0.338	−23.2	57	Blue, nuclear. Visible for ∼300 days.
2020scb	SN Ic	⋯	0.018	−17.7	56	Visible for ∼130 days.
2021zj	SN II	⋯	0.046	−18.6	56	Flash ionization followed by CSM interaction.[h]
2020acty	SN Ib	⋯	0.047	−17.9	55	Visible for ∼70 days.
2018lnb	SLSN-II	⋯	0.222	−22.0	52	Incorrect host association. Visible for ∼175 days.
2021axu	TDE	⋯	0.192	−21.4	51	Incorrect host association.[i] Visible for ∼160 days.

Notes. The bold text designates a transient event that is likely anomalous, and we consider successfully tagged. We consider 26/29 (90%) objects as anomalies due to their (updated) classification, or peculiar light-curve behavior and/or host galaxy, as opposed to the 15/29 originally considered anomalous based solely on the original TNS classification labels.

[a] The updated classification was made by Yao (2023) after our training process.

[b] Spectra displays narrow H_α at early times, then normal width at late times with broadened Lorentzian H_α profile (no clear P-cygni). Coupled with slow rise time and bright peak, we favor a IIn classification.

[c] CSM interaction has been posited as the likely driver for the secondary bump seen in this light curve (Soraisam et al. 2022).

[d] See Quimby et al. (2018) and Gal-Yam (2019) for details.

[e] See Munoz-Arancibia et al. (2023a) for details.

[f] ≈1.4% core-collapse SNe. See Sit et al. (2022) for details.

[g] Frederick et al. (2021) report this object to be a Trakhtenbrot AGN flare (Trakhtenbrot et al. 2019).

[h] See Jacobson-Galán et al. (2024).

[i] Association is correct in current `GHOST` version.

statistics for the anomalous classes, we cannot say that our algorithm will detect 75% of all TDE, for example.

When considering other distributions, such as redshift, we find that our anomaly detection algorithm (which does not use redshift as input) is agnostic. Unlike the IF model, our RFC model tags objects that encompass the entire redshift range of ZTF. We find $z = [0.011, 0.338]$, a mean redshift of $\bar{z} = 0.087$, a median redshift of $z_{\rm med} = 0.052$, and a standard deviation of $\sigma_z = 0.079$. The bounds and standard deviation of the tagged redshift distribution imply that our model has learned to preferentially tag intrinsically bright (correlated with high redshift due to Malmquist bias, which favor SLSNe, TDEs) and intrinsically dim (overrepresented in nearby, low-redshift galaxies, which favor SESNe) events. Meanwhile, the median redshift aligns with the median redshift of ZTF-BTS (see Figure 4 of Fremling et al. 2020), which is correlated with the redshifts of the majority of "Normal" objects. Using redshift measurements, we can transform apparent magnitudes into absolute magnitudes using a flat ΛCDM cosmology with $H_0 = 70\ {\rm km\ s^{-1}\ Mpc^{-1}}$, and $\Omega_M = 0.3$. We find that the tagged objects span the gamut of extragalactic transient peak absolute magnitudes: $M \in [-23.2, -15.2]$ mag, evidence that our model is able to find objects at all intrinsic brightness scales (Richardson et al. 2014).

3.1.2. Spectroscopic, Contextual, and Behavioral Anomalies (After Vetting)

As we describe at the beginning of Section 3, we consider a transient to be an anomaly if it falls into any one (or more) of the three categories: spectroscopic anomaly, contextual anomaly, and behavioral anomaly. The former can be obtained via classification of the SN spectrum and without careful inspection of the light curve or host galaxy, except in the case of reclassification. Additionally, the latter two require such expert vetting to identify peculiar behavior (which can occur, e.g., after the spectrum was taken) or host galaxy environment. When we consider the results after expert vetting, our purity is greater than what we were led to believe from the TNS spectroscopic label alone. Objects that fall into any one or more of these three categories after careful manual vetting by an expert post facto are in bold text in Table 4. We consider all such objects a successfully tagged event by our model. Those without bold text are likely to be nonanomalous, and thus represent a "misclassification." We offer several insights into common failings of our model in Section 6.

For completeness, we carefully inspect the light curves, spectra, and host galaxies of all 455 test set objects. This was done consulting photometry, spectra, host galaxy postage stamps, redshift estimates, catalog information, and other resources available on the TNS, ANTARES (Matheson et al. 2021), and ALeRCE (Förster et al. 2021) brokers, and YSE's transient survey management platform `YSE-PZ` (Coulter et al. 2022). We find 16 transients originally categorized as "Normal" that are updated to the anomalous "Anomaly" classification, resulting in 68 total and a 31% overall increase in anomalies. Of these 16 new anomalies, 11 (or ∼69%) were presently tagged by our model as anomalous, of which we argue eight (SN 2021uzt, SN 2020aeuh, SN 2021fyp, SN 2021cpi, SN 2021ackd, SN 2021ckb, SN 2021adpx, and SN 2020ywo) were previously misclassified on TNS (and therefore, should be considered spectroscopic anomalies), and three can be characterized as behavioral anomalies: two (SN 2020abah, SN 2020rmk) as members of the rare, long-rising (>40 days) SN II class (≈1.4% CC SNe; see Sit et al. 2022 for details), and one (SN 2021zj), with two peaks separated by 100 days, exhibits flash ionization followed by CSM interaction (Jacobson-Galán et al. 2024). Note that all reclassifications were made with the original classification spectrum uploaded to TNS; we did not use any additional proprietary spectra. It follows that there remain five objects that were incorrectly not tagged as anomalous but appear to be of an anomalous nature. Specifically, SN 2021wun is also a known member of the long-rising SN II class (Sit et al. 2022); SN 2020eyj has previously been described as having late-time CSM interaction (Fremling et al. 2020); SN 2021yfi has a peak absolute magnitude of −21.3 mag (but lacks narrow emission lines, which prevented a superluminous IIn classification[46]); SN 2019bcv is uncommonly bright (∼−19.5 mag), red ($g-r \approx 1$ mag), and visible for ∼600 days; and SN 2021ttg is reclassified from SN Ia to SN Ia-91bg-like in this work, prompted by an ANNs search and revetting of the spectrum (see Table 8 and Section 5.2).

Now, post facto, we find 26/29 objects successfully tagged anomalous, resulting in an effective purity of ∼90%. Note that, of the three remaining objects considered "Normal," one (33%) has an incorrect host association (SN 2021aff), which is a possible explanation for its misclassification.[47] The second (SN 2020tnq) has no observations until approximately +26 days after peak (estimated from the evolution of the spectrum), possibly tricking our model into inferring the object is intrinsically fainter (and thus more likely to be of anomalous SESNe classes) than it is. The third (SN 2021our) is a normal SN Ia from the spectrum in a standard elliptical galaxy, but is underluminous based on peak absolute magnitude calculations ($M_{\rm abs} \sim -18.3$ mag), the strongest evidence for the misclassification.

After vetting and reclassifications, we update the per-type percentage of objects tagged by the anomaly detection (AD) model according to the updated spectroscopic label. The tagged objects of "Normal" spectroscopic classes drops from nine SN II to five SN II (7.3% of all SN II in test set) and from five SN Ia to two SN Ia (0.06% of all SN Ia in test set). Similarly, for our anomalous classes, our model tagged seven TDE (77.8%), three SN IIn (26.7%), two SLSN-II (66.7%), two SN IIb (33.3%), one SN Ib (16.7%), one SN Ic (11.1%), and one SN Ic-BL (20.0%). Our model additionally tags 2/3 SN Ia-91bg-like (66.7%) and 2/3 SN Ia-CSM (66.7%), both classes that were identified through reclassification and not included originally in our training set.

Subsequently, we show updated confusion matrices in Figure 3, which describe the effectiveness of our model at classifying *spectroscopic*, *contextual*, and *behavioral* anomalies within the observed ZTF Alert Stream. Now, we record an accuracy of 90%. We score a "Normal" completeness of 99% at a "Normal" purity of 90% while achieving an "Anomaly" completeness of 38% at an "Anomaly" purity of 90%. As designed, our algorithm indeed prioritizes anomaly purity over sample recall, and is effective at identifying SNe with peculiar attributes in regards to the spectroscopic label and beyond.

3.1.3. Additional Performance Evaluation

Beyond a confusion matrix, there are other methods to evaluate the performance of a binary classifier. With an RFC, the final prediction is an aggregate of the final prediction of each tree in the forest, and so, the fraction of votes belonging to either "Normal" or "Anomaly" can be conceptualized as a classification probability. Traditionally, the final classification is assigned as a simple majority, where 50% is the decision threshold, but this choice can be changed to reflect different aspects of the model. It follows that a higher decision threshold considers only events with near-unanimous decision by the trees in the forest, whereas a lower decision threshold is the opposite case; such a trade-off is construed by a receiver operating characteristic (ROC) curve. The ROC curve measures the rate of true positives and false positives as a function of the decision threshold from 0 to 1, and the model accuracy represents a single point along the curve. Meanwhile, the area under the curve (AUC) quantifies the separability of our two classes, where in the limit of perfect classification the AUC approaches unity. A high AUC indicates a high true positive rate and a low false positive rate, whereas a low AUC indicates a low true positive rate and a high false positive rate.

In the left panel of Figure 4, we show the ROC curve for our model, reporting the AUC for only the anomalous "Anomaly" class in the cases of random guessing (red), *spectroscopic* anomalies only before vetting (blue), and any of *spectroscopic*, *contextual*, *behavioral* anomalies after vetting (green). We also report the AUC standard deviation across our 25 different random seed iterations (shown as faded lines, with the bold line denoting our final model). We find by the spectroscopic label alone before vetting we achieve an AUC of 76% ± 1%, and for all anomaly categories after vetting, we achieve an AUC of 83% ± 1%. The greatest separation between these two anomaly criteria occurs at the low true positive and low false positive rate regime, which indicates that the model's

---

[46] See classification report from https://www.wis-tns.org/object/2021yfi.

[47] We remind the reader that the host association is performed upstream, and as in the real-time data processing scenario, an expert cannot a priori know the correctness of the association of a flagged anomaly before human-on-the-loop vetting. Erroneous associations (estimated to be ∼5% by Gagliano et al. 2021) can be identified and dismissed by an expert, underscoring the indispensable role of human intervention in the analysis pipeline.

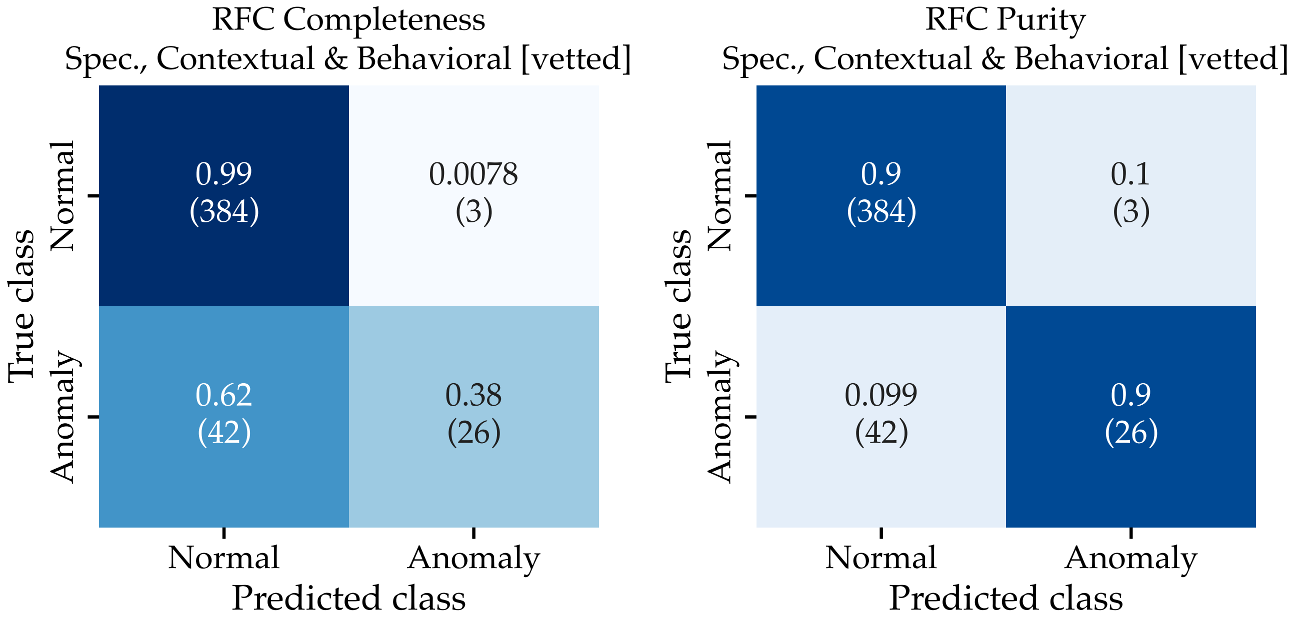

Figure 3. Same as Figure 2, but including additional *spectroscopic*, *contextual*, and *behavioral* anomalies after expert vetting of the classification spectra, host galaxy environment, and light-curve evolution, respectively. We achieve 90% effective purity of detecting *spectroscopic*, *contextual*, and *behavioral* anomalies combined. We find this effective purity score is consistent with the spectroscopic subset of transients identified as anomalous in our real-time deployment on the nightly ZTF Alert Stream via ANTARES (∼83%, see Table 5).

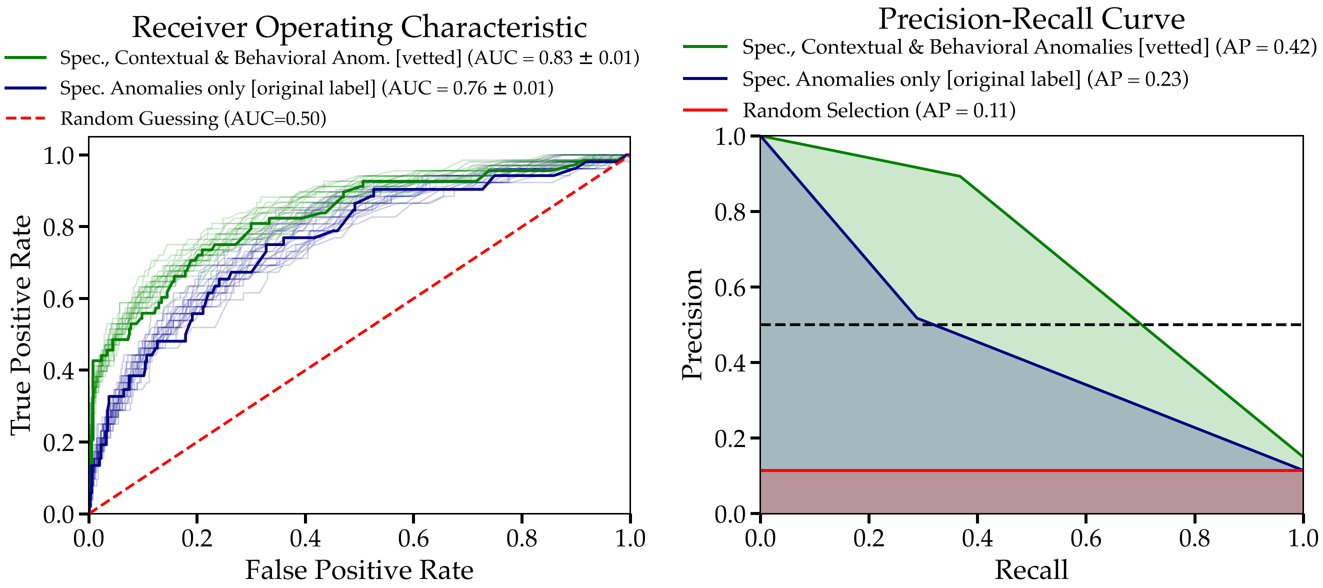

Figure 4. Left: receiver operator characteristic (ROC) curves for *spectroscopic*, *contextual*, and *behavioral* anomalies after vetting (green solid line), *spectroscopic* anomalies only before vetting (blue solid line), and by random selection (red dashed line). Each curve is a different initialization from 1 of 25 different random seeds, where the bold line is the chosen model's performance. The area under the curve (AUC) for the "Anomaly" classification is listed along with its standard deviation across the 25 different random seeds. Right: precision–recall curves for *spectroscopic*, *contextual*, and *behavioral* anomalies after vetting (green solid line), *spectroscopic* anomalies only before vetting (blue solid line), and by random selection (red solid line), where any object with $P(\text{anom}) < 50\%$ is cast to 0, and $P(\text{anom}) \geqslant 50\%$ is cast to 1. The average precision (AP) score is shown in parentheses, and a 50% purity is shown as a dashed black line. Our anomaly detection model significantly outperforms random selection, and can achieve high levels of purity ($\gtrsim$50%) at a range of recall thresholds ($\lesssim$70% after vetting, $\lesssim$32% before vetting).

improvement primarily benefits in the detection of true positives (anomalies) while maintaining low false positives when vetted anomalies are considered.

In the right panel of Figure 4, we show a precision–recall curve for our model performance. We consider any object with $P(\text{anom}) < 50\%$ as "Normal" and $P(\text{anom}) \geqslant 50\%$ as "Anomaly." At a $P(\text{anom}) = 50\%$ threshold, we achieve a purity of *spectroscopic* anomalies of approximately 50% (52%), which optimizes the trade-off between precision and recall for our use-case—we achieve the maximum recall for which we identify more anomalies than nonanomalies. For this threshold visualized in our precision–recall curve, we cast any object with $P(\text{anom}) < 50\%$ to $P(\text{anom}) = 0\%$ and any object with $P(\text{anom}) \geqslant 50\%$ to $P(\text{anom}) = 100\%$. This manifests as an inflection point, occurring at the "Anomaly" recall values as shown in the confusion matrices (29% before vetting, 38% after vetting). We show the average precision[48] from prediction scores in parentheses, denoted as the

---

[48] https://scikit-learn.org/stable/modules/generated/sklearn.metrics.average_precision_score.html

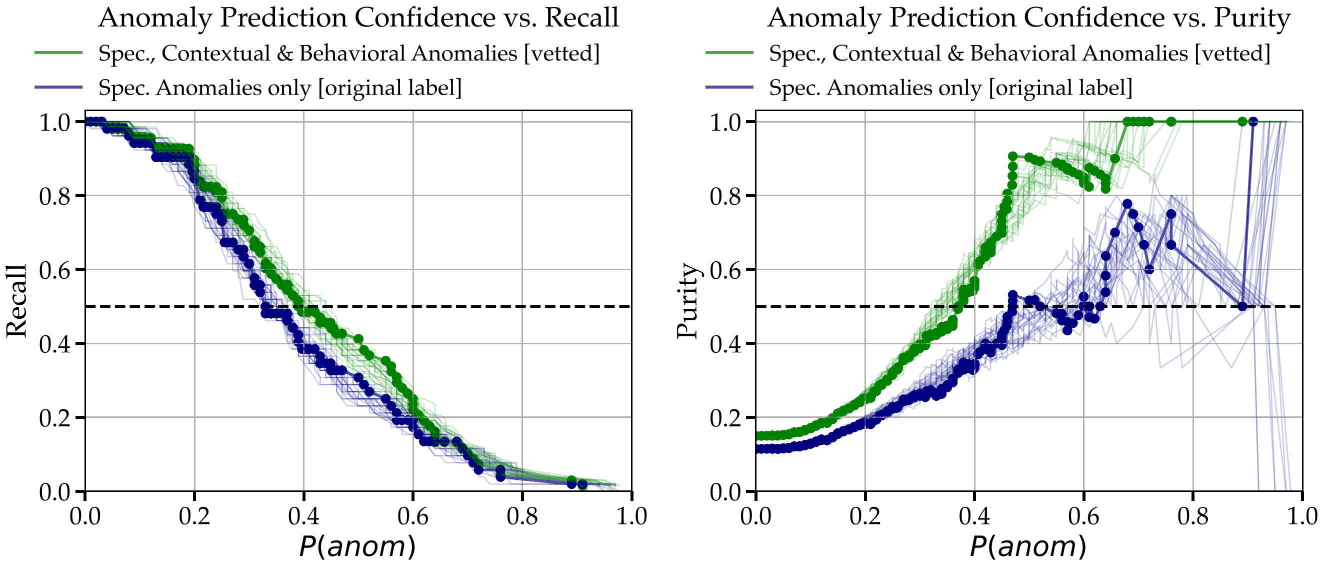

Figure 5. The observed completeness (left panel) and observed purity (right panel) of our model as a function of anomaly classification confidence, P(anom), calculated at *full* light-curve phase. Individual object confidence scores are represented as circles. Each curve is a different initialization from 1 of 25 different random seeds, where the bold line is the chosen model's performance. At our chosen thresholds of $P(\text{anom}) = 50\%$ for this work, we achieve an overall purity score of 52% and a completeness of 29% for *spectroscopic anomalies* only before vetting (blue) and an overall purity score of 90% and a completeness of 38% for *spectroscopic*, *contextual*, and *behavioral* anomalies after vetting (green), reflected in the confusion matrices of Figures 2, 3, respectively.

weighted mean of precision values achieved at each threshold, using the increase in the recall from the previous threshold as the weight. Our anomaly detection model significantly outperforms random selection, and can achieve high levels of purity ($\geqslant$50%) at a range of recall thresholds: $\leqslant$32% for *spectroscopic* anomalies only before vetting, and $\leqslant$70% for vetted *spectroscopic*, *contextual*, and *behavioral* anomalies combined. For even higher levels of purity ($\geqslant$75%), these recall thresholds are $\leqslant$15% for *spectroscopic* anomalies only before vetting, and $\leqslant$49% for vetted *spectroscopic*, *contextual*, and *behavioral* anomalies combined.

Ultimately, we want to better understand on a granular scale how reliably we can trust the model output at varying confidence scores, particularly to investigate if higher confidence correlates in the increased likelihood of an object being anomalous. In Figure 5, we show the observed completeness (left panel) and observed purity (right panel) of our model as a function of anomaly classification confidence, P(Anomaly) (hereafter P(anom)). Because of our hyperparameter selection, there are fewer high confidence scores ($P(\text{anom}) \geqslant 50\%$) than low ones ($P(\text{anom}) < 50\%$), and of those that have $P(\text{anom}) \geqslant 50\%$, the majority are grouped within the 50%–70% range, with only a few greater than 70% (see column (6) of Table 4).

In our test set before vetting, every object with $P(\text{anom}) \geqslant 70\%$ is anomalous except one, SN 2021uzt, which has the overall second highest anomaly score at $P(\text{anom}) = 89\%$. A possible driver of the high anomaly score is the likely incorrect host association for this object (the real host is not visible). SN 2021uzt was originally classified as an SN IIn (Tucker 2021), but later classified as an SN II (Chu et al. 2021). Due to this, SN 2021uzt is represented as the sharp dip shown in the right panel in blue, and is the sole reason for the decline in purity at high confidence scores. However, upon vetting and reevaluating the two spectra, there is evidence of a IIn-like Lorentzian H_α profile and none of P-cygni. Moreover, SN 2021uzt is bright at its peak ($M_{\text{abs}} \sim -19.4$ mag) with a long rise (>30 days). This evidence suggests the original IIn classification is the best characterization, and we adopt it for this work. Thus, after vetting (green), all objects with $P(\text{anom}) \geqslant 70\%$ are anomalous.

Averaged over the 25 different seed iterations, the purity of vetted anomalies steadily increases with rising confidence score from 0% to a peak around $P(\text{anom}) \approx 45\%$, plateaus until $P(\text{anom}) \approx 65\%$, then rises to a perfect purity at $P(\text{anom}) \approx 70\%$ and sustains it until a maximum classification confidence.[49] If $P(\text{anom}) \geqslant 47\%$, the purity achieved is $\geqslant$80% (the nearby $P(\text{anom}) = 50\%$ is reflected in the purity confusion matrix of Figure 3). This trend shares similarities when considering only *spectroscopic* anomalies (before vetting) except the difference that stems from SN 2021uzt and that the maximum purity that can be achieved is overall lower; if $P(\text{anom}) \approx 47\%$, the purity achieved before vetting is nearly half that compared to vetted anomaly candidates, achieving ∼40% at worst (but usually $\geqslant$50%). Perhaps the most interesting insight gleaned from Figure 5 is that if we vet our anomaly candidates, we achieve a 50% purity score at a lower P(anom) threshold (38%) than if we do not (47%). Fluctuations in the model from random seed iterations impact the outcome of the anomalies before vetting more severely than after vetting, likely due the fact that more objects with high anomaly scores are deemed anomalous after vetting than vice versa.

In a similar manner, the recall steadily decreases with increasing confidence score across the range of recall scores, alluding to the fact that the majority of objects have low anomaly confidence scores (because most are truly nonanomalous), and a minority have high anomaly scores. This behavior reflects the distribution of anomalies we observe. There is little difference in the recall score before and after vetting at the P(anom) margins; the greatest gain (≈10% in recall) comes around $P(\text{anom}) \sim 50\%$, although we observe that this large

---

[49] As we will see in Table 7, $P(\text{anom}) \geqslant 70\%$ does not guarantee perfect purity; however, in the case of Table 7, our model is not trained on features extracted from YSE photometry, and we consider the score at *any* light-curve phase instead of at the end of the full phase.

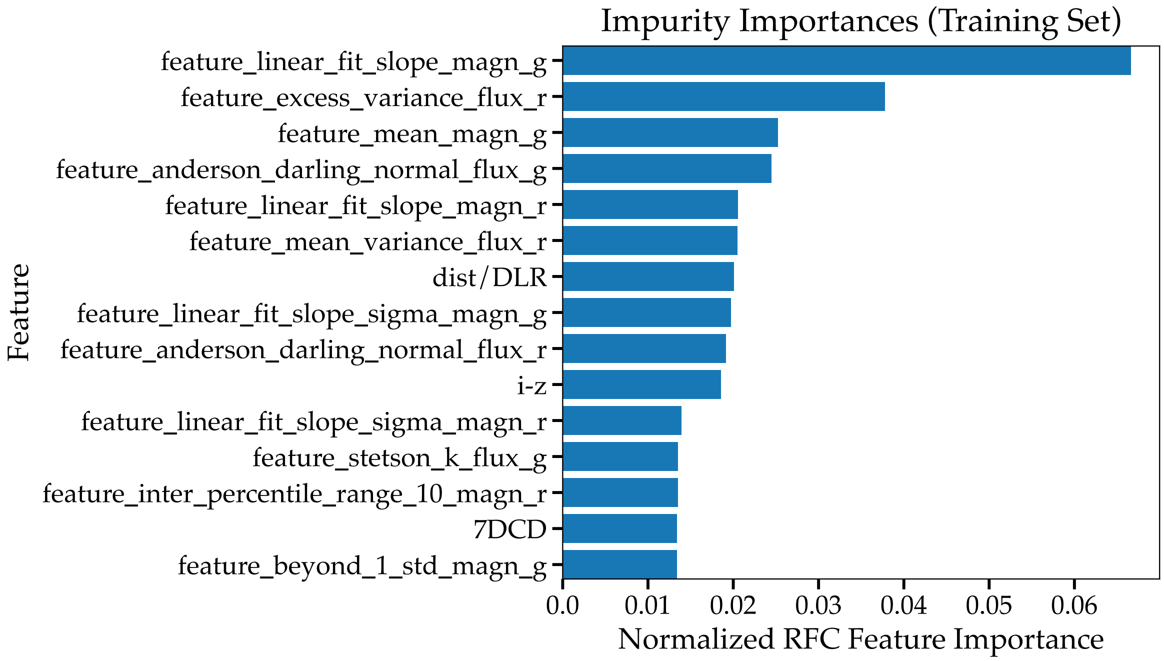

Figure 6. The top 15 light-curve and host galaxy features, ordered by the greatest normalized impurity-based feature importance in our binary "Normal" vs. "Anomaly" task as determined by our final random forest classifier model. The features with highest importance can be broadly categorized by light-curve slope fit and error, *g*-band brightness and its variance, transient radial separation, and host galaxy color. The high importance of such host galaxy features is consistent with findings from Gomez et al. (2020) and Gagliano et al. (2021) despite their different classification tasks (SLSNe vs. non-SLSNe, and SNe Ia vs. CC SNe, respectively).

offset is in part attributed to fluctuations from the random seeding.

For our chosen model, we achieve a maximum purity of 100% at a recall of 13% given a confidence score of P(anom) = 68%. But the best balance achieved between purity and recall is at P(anom) = 50%, which is reflected throughout this work.

3.2. Feature Importance

To better understand our model, we desire to determine which light-curve and host galaxy features (see the complete list in Appendix A) are most valuable for our anomaly detection task. A common method is to calculate impurity-based feature importances, computed as the (normalized) total reduction of the criterion brought by that feature (Gini importance). We use the `feature_importances_` attribute from `sklearn.ensemble.RandomForestClassifier`, and display the results in Figure 6. As is standard procedure, we use the normalized feature importance, defined as the percentage of times the feature is used as a split, as our metric for feature significance.

Under this method, we find that the most important light-curve features are the light-curve slope in a least squares fit of the linear stochastic model with Gaussian noise described by observation errors $\{\delta_i\}$ (`feature_linear_fit_slope_magn_{g,r}`) and its error in g band (`feature_linear_fit_slope_sigma_magn_g`), the mean and excess variance r-band flux (`feature_excess_variance_flux_r`, `feature_mean_variance_flux_r`), mean magnitude in g band (`feature_mean_magn_g`), and the unbiased Anderson–Darling normality test statistic for flux (`feature_anderson_darling_normal_flux_{g, r}`). We suspect that the `feature_linear_fit_slope_magn_g` at 0.067 is the most important feature overall at nearly twice the contribution of the second most important feature (0.038; `feature_excess_variance_flux_r`) because many rarer classes of transients like TDEs show consistently strong blue colors or SESNe, which tend to exhibit constantly weak blue colors relative to red throughout the light-curve evolution.

Interestingly, like FLEET (Gomez et al. 2020, 2023a, 2023d), which is designed to find SLSNe and TDEs with an RFC, we find the most important host galaxy feature is the projected angular separation between the transient and host galaxy normalized by host radius (`dist/DLR`, denoted as θ/d_{DLR} in the original `GHOST` paper; Gagliano et al. 2021). This is an indication that SN classes, particularly rare ones, may preferentially occur at different locations throughout their host galaxies than more normal SNe. For example, TDEs are nuclear events, and would predominantly have a small `dist/DLR` value, making this potentially a powerful indicator for separating TDEs from non-TDEs.

The second most important host feature is the aperture magnitude $i - z$ color difference (`i-z`), followed by `4DCD`, the four-dimensional color distance in $g - r$, $r - i$, $i - z$, and $z - y$ from the PS1 stellar locus (Tonry et al. 2012). This is the path traced by stars in color–color space (see Section 2.2 of Gagliano et al. 2021 for details). Color-derived features encode information about the metallicity, mass, and SFR of host galaxies. These characteristics are known to be associated with the type of SN (e.g., Hansson et al. 2012), including host galaxy property correlations with SNe Ia (e.g., Johansson et al. 2013; Henne et al. 2017; Kelsey et al. 2023). Furthermore, the color of galaxies can effectively distinguish between early- and late-type galaxies (Strateva et al. 2001; Nair & Abraham 2010). This suggests that valuable host galaxy features for classifying SNe, and potentially anomalous SNe, are related to previously established galaxy correlations.

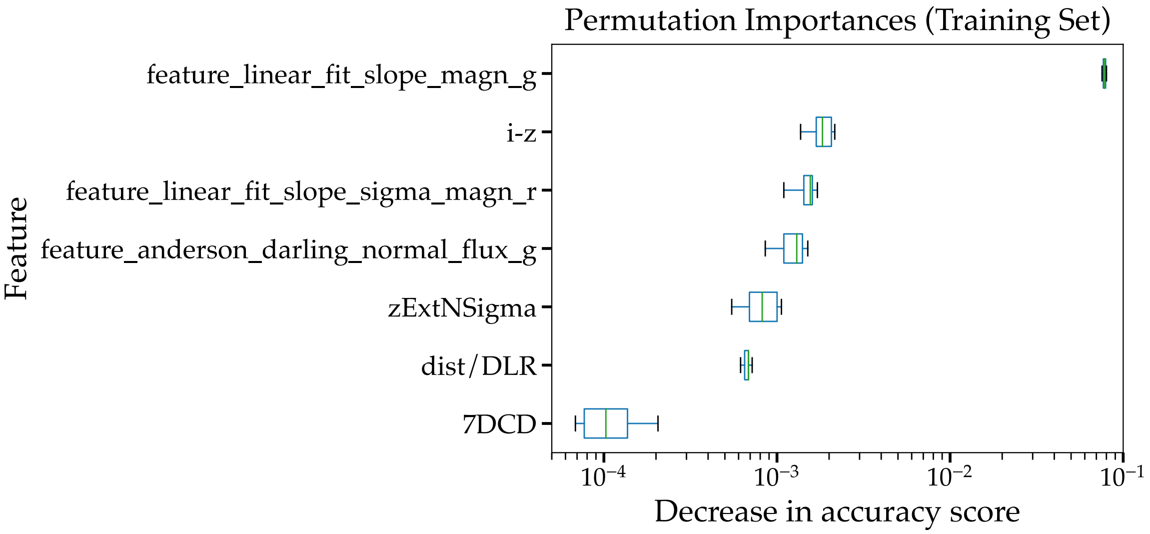

Figure 7. A boxplot of the top seven most important features from our training set with the highest permutation importance using the method from Breiman (2001). We use 10 permutations per feature, finding the features that most negatively impact accuracy. All features except `zExtNSigma` are shown as among the most important identified by the impurity-based feature importance metric (see Figure 6). The light-curve feature `feature_linear_fit_slope_magn_g` remains the singular most important feature for identifying anomalies for this work.

Moreover, our host galaxy feature importance results align very well with those of Gagliano et al. (2021) when considering radial offset and color-derived features. They used a gradient boosting model to classify SN Ia versus CC SNe from host galaxy information alone derived from `GHOST` host galaxy associations, finding their 10 most important features for classification were of the same three main categories: radial offset, including θ and θ/d_{DLR} (their #1 and #2 most important features, respectively); color-derived features, including `4DCD`, `g-r`, `g-i`, `r-i`, and `i-z` (their #7, #3, #4, #5, and #9 most important features, respectively); and morphological features, including `momentXX` in g and i and `ExtNSigma` in g (their #10, #8, and #6 most important features, respectively).

In this work, we do not use θ, `g-r`, `g-i`, `r-i`, but we do use `dist/DLR` (θ/d_{DLR}), `i-z`, `4DCD`, `{g,r}momentXX`, and `gExtNSigma`. For us, `dist/DLR` is our most important host galaxy feature, and `4DCD` is our second most important, with `rmomentXX`, and `gExtNSigma` being in the top half of host galaxy features, and `gmomentXX` being in the lower half of important host features. Thus, even though our classification task is different from that of Gagliano et al. (2021), there is great overlap in the most important features chosen in regards to radial offset (`dist/DLR`) and color-derived features (`i-z`, `4DCD`). Further exploration and quantization of the usefulness of these features are areas of future work.

However, it is known that impurity-based feature importances can be misleading for high cardinality features and artificially inflate numerical feature importance.[50] Because of the strong correlations between the majority of host galaxy features, the variable importances returned by our RFC model may not reveal the most valuable features for SN classification, or paint the entire picture. Thus, we opt to additionally investigate feature importance using the permutation importance method described in Breiman (2001). Permutation importance measures the impact of each feature on the model's performance by evaluating how much the model's accuracy (or another evaluation metric) degrades when the values of that feature are randomly shuffled. Features that have a high impact on the model's accuracy will result in a significant decrease in accuracy when permuted. Therefore, it provides a measure of the feature's importance in making accurate predictions. We display the top seven features with greatest permutation importance in Figure 7.

Six of the seven top features identified by permutation importance align well with those identified by impurity-based importance. Specifically, `feature_linear_fit_slope_magn_g` remains the most important by a wide margin over every other feature, indicating that particular g-band slope profiles are the best indicator of anomalous activity. Light-curve features `feature_linear_fit_slope_sigma_magn_r` and `feature_anderson_darling_normal_flux_g` remain as among the most important. As before, radial offset (`dist/DLR`) and color-derived features (`i-z`, `4DCD`) for host galaxy features are vital. Meanwhile, the only feature flagged as important by permutation and not by impurity methods is `zExtNSigma`. Because of our highly dimensional feature-space, the individual impact of any one feature on the final anomaly classification (besides these top few) is small.

3.3. Filter Deployment, Real-time Anomaly Detection, and Follow-up

After demonstrating the success of our anomaly detection model, we deploy it in a real-time scenario: the ZTF Alert Stream. We compose a *Filter* on the ANTARES broker, which first requires objects to have been tagged successfully upstream by the `lc_feature_extractor` *Filter*. Then, we perform the same preprocessing cuts to initially create our databank as outlined in Section 2.2, except that we do not query PS1-PSC to save on computational time in a real-time scenario. We then apply our trained RFC model. Calculated light-curve features,

---

[50] See https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html#sklearn.ensemble.RandomForestClassifier.feature_importances_.

host galaxy features, and anomaly score are saved to the tagged loci as locus properties.

All objects that successfully pass through the *Filter* are tagged as "`LAISS_RFC_AD_filter`" regardless of anomaly score. However, those with $P(\text{anom}) \geqslant 50\%$ (or any threshold of a user's choosing) at any point during their light-curve evolution are processed by our Slack API App[51] as a bot that sends out a notification of the IAU name, ANTARES Locus ID, TNS spectroscopic classification (if it exists), and the current anomaly score. An expert then manually inspects each anomaly candidate. Our bot queries the ANTARES broker using the `antares_client` package (NOIRLab CSDC Team 2023), and collates all objects tagged by `LAISS_RFC_AD_filter` within any number of days of the expert's choosing. From there, we apply Wide-field Infrared Survey Explorer (WISE; Wright et al. 2010) color selection criteria for AGNs from Jarrett et al. (2011) and Stern et al. (2012) to remove likely AGN candidates with high anomaly scores. We typically run the bot once a day to vet the previous nights' up-to-date and active anomaly candidates.

Our *Filter* was deployed on 2023 August 22. The data cutoff for this work for our real-time deployment is 2023 October 22, 2 months since deployment. In that short time, `LAISS` has successfully processed ∼1200 loci, of which 45 have achieved an anomalous classification of $P(\text{anom}) \geqslant 50\%$ and are listed in Table 5, ordered by maximum anomaly score, $\max(P(\text{anom}))$. In total, 39 are transients (21 with spectroscopic classifications), and 6 are AGN or AGN candidates (1 spectroscopically confirmed).

For the spectroscopic sample, we consider 17/21 (81%) to be *spectroscopic*, *contextual*, or *behavioral* anomalies (which aligns closely to the 86% purity of our spectroscopic test set), marked by bold text. Note that the objects here are a mix of those that were already classified and active at the time of *Filter* deployment (e.g., TDE 2022fpx), those that were tagged by our model as anomalous before a classification spectrum was acquired by others (e.g., SN Ib 2023nlj), and those whose spectra were acquired by us—indicated by a dagger symbol (e.g., SN Ia-CSM 2023ocx). We report new and/or updated classifications from our acquired spectra to TNS.

The distribution of objects is as follows: six SN IIn, three SN Ia-CSM, one M31 Nova, two TDE, one SLSN-II, one SN Ib, one SN IIb, one SN Ia-91bg-like, three SN II, and two SN Ia. This is further evidence that our algorithm is capable of identifying an array of different spectroscopic anomalies with high purity in practice, with a preference at identifying long-lived CSM-interacting events like SNe IIn and SNe Ia-CSM. Moreover, as we identified in the spectroscopic test set, our algorithm tags anomaly candidates at all common redshift ranges observable by ZTF: $z \in [0.029, 0.194]$. Likewise, these objects are varied in their apparent and intrinsic brightness, host galaxy type, morphology, and apparent size, with no obvious difference compared to the test set.

Of our nine obtained spectra, we consider six to be anomalous. In fact, one was used to aid in the reclassification of a normal Type II to a IIn, as in the case of SN 2023nof (Aleo et al. 2023). The three remaining are of a normal Type Ia and two Type II SNe, although two were only (marginally) above the anomaly score threshold for one epoch in its light-curve evolution, and the other had an incorrect host association that may have played into the misclassification. Generally, we find throughout this work that objects with higher anomaly scores and/or are above the $P(\text{anom}) = 50\%$ anomaly threshold for more epochs are more likely to be anomalous, supported in part by the evidence from our purity as a function of prediction confidence as shown in Figure 5.

Three objects have likely incorrect host associations from `GHOST` (SN 2023omf, SN 2023khp, SN 2023otw), although oddly two of three are of anomalous classes. This could indicate that, in some cases, truly anomalous phenomena as captured by the light-curve evolution alone (or the pairing of the light-curve evolution with an incorrect host that happens to be atypical for that light curve) is sufficient to be flagged anomalous. Although we make no formal analysis of misclassified hosts, Gagliano et al. (2021) estimates that the misassociation rate of `GHOST` is approximately 5%. Continual improvements to the pipeline and addition of catalogs such as GLADE (Dálya et al. 2018) since the original release likely will decrease the rate of misassociations.

Perhaps the most promising revelation from our real-time deployment beyond the high purity is that occasionally our algorithm flags a transient as anomalous *well before* ($\gtrsim$2 weeks) the first classification spectrum. When we reextract the light-curve features for all epochs of all tagged transients and reapply our model throughout its light-curve evolution as a mock up for real-time deployment, we can investigate when these candidates *would have been tagged* had our model been deployed before the transient was active.

Of our spectroscopic subset, two (TDE 2022fpx, SN Ib 2023nlj) were tagged 39 and 16 days before their respective classification spectra. With earlier anomaly classifications, we can more closely study the physics linked to its explosion (Gagliano et al. 2023), as well as more completely follow up the most interesting events with spectrophotometric resources and create a fuller picture of the spectral energy distribution (SED) evolution (Pierel et al. 2018; Vincenzi et al. 2019). SED templates enable the construction of models and simulations that better capture genuine anomalies, supplementing existing idealized models focused on typical phenomena. Realistic simulations of anomalies can enhance machine learning algorithm training in classification and anomaly detection tasks—contrasting with the tendency of algorithms trained on idealized simulations to mischaracterize anomalies (see, e.g., Muthukrishna et al. 2022; Aleo et al. 2023).

For the photometric sample, unfortunately, there is no way to confirm the anomalous nature of 18 anomaly candidates. However, we have proffered potential classifications and anomaly candidates for completeness. Potential classifications in Tables 4–7 come from a combination of FLEET (Gomez et al. 2020) predictions, the YSE DR1 ParSNIP and SuperRAENN classifiers (Aleo et al. 2023; adapted from Boone 2021; Villar et al. 2020, respectively), SN Ia SALT3 fits (Kenworthy et al. 2021), and an expert analysis. Although we cannot use the photometric subset to quantify our algorithm's results, we make an inference at their anomalous nature by highlighting in bold text those that are strong anomaly candidates (as we do with the spectroscopic subset).

Despite the lack of a classification spectrum, some objects stand out as likely anomaly candidates. For instance, AT 2021rjf, which exhibited no previous variability, shows signs of a normal SN II light curve followed by a second bump likely powered by CSM interaction that has lasted 800 days and

---

[51] https://api.slack.com

Table 5

Real-time Search: *Transient-only* Loci with $\max(P(\mathrm{anom})) \geqslant 50\%$ at *Any* Light-curve Phase

45 Total Tagged Loci, `LAISS_RFC_AD_filter` Filter, Ordered by $\max(P(\mathrm{anom}))$

ZTF ID	IAU Name	Spec. Class	z	$\max(P(\mathrm{anom}))$	Remarks
Spectroscopic					
ZTF22aadesap	2022fpx	TDE	0.073	0.92	Tagged anomalous 39 days before SGLF’s spectrum.
ZTF23aatzhso	2023oom	Nova	⋯	0.87	In M31.
ZTF23aaspcfl	2023nlj	SN Ib	0.03	0.83	Tagged anomalous 16 days before ZTF’s spectrum.
ZTF23aanptpp	2023koq	SLSN-II	0.104	0.77	Tagged anomalous at peak.
ZTF23aarktow	2023myo	SN IIb[a]	0.035	0.74	Reclassified. First ZTF-r epoch is shock cooling.
ZTF22aaetqzk	2022gzi	SN IIn	0.089	0.72	⋯
ZTF23aamsetv[b]	2023kvy	TDE	0.16	0.67	Tagged on minimum required ZTF-g, r obs.
ZTF23abcinum[c]	2023sds	SN II	0.065	0.67	YSE target. In elliptical galaxy?
ZTF23aajkisd	2023iex	SN IIn	0.029	0.66	⋯
ZTF23aalgqsq	2023jdh	SN IIn	0.054	0.65	Faint host.
ZTF23aatdcey[c]	2023nof	SN IIn	0.069	0.65	With our spectrum, reclassify from Type II to IIn.
ZTF23aapgswu	2023mcs	SN Ia	0.03	0.64	First epochs are ≈+20 days after peak.
ZTF23aatcsou[c]	2023nwe	SN IIn	0.194	0.64	Tagged anomalous 14 days before our spectrum.
ZTF23aagpjyp	2023ggb	SN Ia-CSM	0.08	0.59	⋯
ZTF23aavtugd[c]	2023omf	SN IIn	0.083	0.57	Incorrect host. Match to SN 1996L.
ZTF23abhafym	2023tsw	SN Ia-91bg-like[d]	0.05	0.57	Reclassified. Red, rapid fading favors 91bg-like.
ZTF23aatabje[c]	2023ocx	SN Ia-CSM	0.076	0.56	Tagged anomalous 22 days before our spectrum.
ZTF23aawblmi[b]	2023otw	SN II	0.087	0.60	Incorrect host.[e]
ZTF23aberpzw[c]	2023swf	SN II	0.024	0.53	Only one anomalous epoch.
ZTF23aamsekn	2023khp	SN Ia-CSM	0.09	0.52	Incorrect host.[f]
ZTF23abayyjm[c]	2023sed	SN Ia	0.14	0.50	Only one anomalous epoch. Spiral host.
Photometric					
ZTF21abiggqx[g]	2021rjf	SN IIn?	⋯	0.80	Long-lived (800 days) CSM interaction?
ZTF23aaveoxd	2023ofr	SN IIn?	⋯	0.76	Faint host. FLEET = 82% SN II.
ZTF23aaqbyzr	2023mic	SN Ia-CSM?	⋯	0.74	YSE target. Second peak likely CSM interaction.
ZTF23abhegfd	2023tjc	SN Ia?	⋯	0.66	FLEET = 59% SN I. SALT3 $c = -0.3$.
ZTF23aaqqeek	2023mne	SLSN?	⋯	0.63	Faint host. FLEET = 82% SLSN-II.
ZTF23abevrtm	2023tim	SN Ia?	⋯	0.63	FLEET = 76% SN I. SALT3 $c = -0.22$.
ZTF23abbbypt	2023sap	SLSN?	⋯	0.61	FLEET = 74% SLSN-II.
ZTF23abaurik	2023she	SN IIP?	⋯	0.60	FLEET = 75% SN II.
ZTF23abekzca	2023tcq	SN Ia?	⋯	0.59	YSE target. FLEET = 85% SN I. SALT3 $c = -0.15$.
ZTF23aazfibd	2023puf	SN II/IIn?	⋯	0.58	Faint host. FLEET = 90% SN II.
ZTF23aasotjh	2023nie	SN II?	⋯	0.57	Peculiar rise. FLEET = 80% SN II.
ZTF23abetluh	2023tdy	SN Ia?	⋯	0.57	FLEET = 46% SN I. SALT3 $c = -0.3$.
ZTF23aaxbkgs	2023pdf	SN II/IIn?	⋯	0.55	Faint host. FLEET = 41% SN II.
ZTF23abdpgvv	2023sws	SN Ia?	⋯	0.55	Incorrect host. SALT3 $c = -0.27$.
ZTF23aaufkak	2023nwk	SN IIn?	⋯	0.54	Faint host. FLEET = 76% SN II.
ZTF23abedgfr	2023syt	SN Ib/c?	⋯	0.54	Poor SALT3 fit. FLEET = 52% SN I.
ZTF23aaewyhm	2023gzn	SN IIn?	⋯	0.50	Behind Sun for ≈150 days. FLEET = 56% SN II.
ZTF23abeujrk	2023tnr	SN Ia?	⋯	0.50	FLEET = 88% SN I. SALT3 $c = -0.29$.

Notes. The bold text designates a transient event that is likely anomalous, and we consider successfully tagged. FLEET does not distinguish between Type II and IIn. Filter tagged one known AGN (ZTF20acvfraq/AGN 2020adpi) and five likely AGN (ZTF22abghche/AT 2023tsa, ZTF23aaloouf, ZTF22abplfmz/AT 2023gld, ZTF23aaxazht, ZTF23aaqxgan/AT 2023tpk), which are not shown in this table.

[a] We reclassify from Type II to IIb.

[b] Spectroscopic follow-up in coordination with FLEET (Gomez et al. 2023d) program.

[c] Targets for which we got spectra and posted their (updated) spectroscopic classification to TNS.

[d] We reclassify from Type Ia to Ia-91bg-like.

[e] Association is correct in current `GHOST` version.

[f] Association is correct in current `GHOST` version.

[g] In the databank used for train/test split (see Table 2).

counting. AT 2023mic also appears to have prolonged CSM interaction. The rest are predominantly SN II/IIn or SLSN candidates in faint host galaxies.

Our `LAISS_RFC_AD_filter` *Filter* continues to run on ANTARES. The full list of processed objects regardless of anomaly score can be found at https://tinyurl.com/LAISSrfcADfilter.

Table 6
Retrospective Search: *Transient-only* Loci with max(P(anom)) ⩾ 50% at *Full* Light-curve Phase

ZTF ID	IAU Name	Spec. Class	z	P(anom)	Remarks
10,000 Randomly Selected Tagged Loci, `iso_forest_anomaly_detection` *Filter*, Ordered by P(anom)					
Spectroscopic					
ZTF22abhwlnm	2022wed	SN IIn	0.114	0.95	Tagged anomalous 118 days before BTDG's spectrum.
ZTF22aadesjc	2022fnl	SN IIn	0.104	0.93	Tagged anomalous 14 days before ZTF's spectrum.
ZTF22abfdzrv	2022vmg	SLSN-I	0.41	0.69	Peak $M_{\mathrm{abs}} \sim -22.4$ mag.
ZTF21abgkfzh[a]	2021qep	SN IIn	0.086	0.67	Faint host. Visible for ∼800 days.
ZTF23aatcsou†	2023nwe	SN IIn	0.194	0.64	Tagged anomalous 14 days before our spectrum.
ZTF23aadjssg	2023dgp	SN Ia-91T-like	0.045	0.64	Likely incorrect host. No rise info.
ZTF18aaiwzie	2023bfv	SN Ia-91T-like[b]	0.086	0.58	Reclassified. No g-band obs for >36 days.
ZTF23aaekebt	2023eqx	SN II	0.02	0.58	Incorrect host?
ZTF23aadbtou	2023cyx	SN Ia	0.033	0.52	No rise info. SALT3 $c = -0.09$.
ZTF22abghrui	2022vwu	SN IIn	0.197	0.50	Tagged anomalous 14 days before ZTF's spectrum.
Photometric					
ZTF21abiggqx[a]	2021rjf	SN IIn?	⋯	0.80	Long-lived (800 day) CSM interaction?
ZTF23aaefpfb	2023fli	SN II?	⋯	0.68	Faint host, incorrect assoc. FLEET≈28% SN II/SLSN-I.
ZTF22aatwxrl	2022oym	SN IIn?	⋯	0.64	FLEET = 51% SN II, 43% SLSN-I. Visible for∼375 days.
ZTF23aaahnss	2023atr	SN IIn?	⋯	0.64	Faint host, incorrect assoc. No rise. FLEET = 60% SN II.
ZTF22absuavp	2022zyh	TDE?	⋯	0.61	Blue. Faint host. Visible for ∼200 days.
ZTF23aafgmaz	2023frg	SN Ib/c?	⋯	0.60	SALT3 $x_1 = +3.00$. FLEET = 84% SN I.
ZTF23aaflptz	2023gbk	SN Ia?	⋯	0.55	No decline. Peak $M_{\mathrm{abs}} \sim -18.3$ mag. Underluminous?
ZTF23aatcola	2023noh	SN Ia?	⋯	0.55	Only one anomalous epoch. FLEET = 53% SN I, 46% SN II.
ZTF23aajestr	2023inr	SN II?	⋯	0.54	YSE Target. Long-rising (∼75 day) SN II candidate.
ZTF23aahjdxa	2023gpp	SLSN/SN IIn?	⋯	0.53	SLSN/IIn? FLEET = 46% SLSN-II, 41% SN II.
ZTF22abrbohu	2022ywi	SN II?	⋯	0.53	Incorrect host? FLEET = 60% SN II.
ZTF23aarzzwu	2023nfs	SN Ia?	⋯	0.52	No decline. Poor SALT3 fit ($x_1 = +3.00$, $c = -0.3$).
ZTF22abzajwl	2023adr	TDE?	⋯	0.51	max(P(anom)) = 0.84. Blue, nuclear. FLEET = 81% TDE.
ZTF23aagxvad	2023glx	SN Ia?	⋯	0.50	Barred spiral host galaxy. No g-band obs after peak.
ZTF23aaempzk	2023fbj	SN Ia?	⋯	0.50	No rise. SALT3 $x_1 = 2.80$, $c = -0.28$. FLEET = 46% SN I.

Notes. The bold text designates a transient event that is likely anomalous, and we consider successfully tagged. We encountered eight AGN or nontransient activity (ZTF19abyfhvp, ZTF21ackmnbo, ZTF21achlwqg, ZTF22aaytzrb, ZTF22aanvqhn, ZTF21abwyelp, ZTF23aaunnwa, ZTF23aaarpdm), which are not shown in this table.

[a] In the databank used for train/test split (see Table 2).

[b] We reclassify from SN Ia to SN Ia-91T-like.

3.4. Retrospective Characterization of the `iso_forest_anomaly_detection` Filter

Our anomaly detection model can also be used for retrospective characterization, similar to that of our spectroscopic test set (Table 4). This enables us to potentially unearth new or overlooked anomalous objects worthy of study or attention. For this assessment, we run `LAISS` on 10,000 random loci tagged (out of over 25,000) by ANTARES' first anomaly detection *Filter*, `iso_forest_anomaly_detection`, and report objects whose maximum anomaly score is ⩾50% at full light-curve phase in Table 6.

The `iso_forest_anomaly_detection` *Filter* is simpler than `LAISS`. It aims to tag transient events like SNe, cataclysmic variables (CVs), and weird or rare events/objects in the night sky. Originally deployed on 2021 April 5, it is an IF algorithm (Liu et al. 2012) trained on 1,000,000 random ANTARES loci using 1000 trees on 106 total light-curve features,[52] 53 for both ZTF-g and ZTF-r passbands. The only selection cuts used enforced that the loci does not reside in the galactic plane ($|b| \geqslant 10°$), and does not have its strongest period between 100 and 1000 days (e.g., to weed out the majority of Mira star contaminants). Thus, it uses no contextual host galaxy information, or selection cuts to reduce AGN, QSOs, or other variable stars. Approximately one out of three tagged objects are on TNS (reported by ANTARES or other teams), and the rest are contaminants. And the vast majority of tagged transients are normal SNe Ia.

Despite the quantity of nontransients and those that are not successfully processed by the `LAISS_RFC_AD_filter` *Filter* (e.g., the locus is a star, the associated host is not found or does not have PS1-*grizy* color information needed for our model, etc.), we do find spectroscopically confirmed and photometric anomalies.

`LAISS` tags 10 objects, of which six fall into the spectroscopic anomaly category only: five SN IIn and 1 SLSN-I. Because one was previously used in the training of our model (SN 2021qep), we exclude this object and say that five of nine are anomalous, or 56% purity. By happenstance, two of these objects (SN 2022vmg and SN 2023nwe) we had confirmed with spectra from previous follow-up campaigns. Of the four spectroscopic nonanomalies, two have likely incorrect host associations (SN 2023dgp, whose potential hosts are vastly different: either a small elliptical or large nearby spiral, and SN 2023eqx, which is between an irregular faint

---

[52] These 106 include all of the light-curve features used in `LAISS` and more—for a comprehensive list, see https://antares.noirlab.edu/properties.

Table 7
Retrospective Search: The 42 Most Anomalous YSE DR1 (Aleo et al. 2023) Events with max(P(anom)) ⩾ 70% at *Any* Light-curve Phase

1153 Processed, from 1975 Total `YSE DR1` Objects, Ordered by max(P(anom))

Spectroscopic (Spec-z)

IAU Name	Spec. Class	z	Peak $M_{\rm abs}$	max(P(anom))	Anom. Obs	Remarks
2020xsy	SLSN-II	0.27	−22.3	0.98	93	Tagged anomalous 13 days before YSE's spectrum.
2020tip	SN Ia-9T-like[a]	0.095	−19.2	0.90	14	Spiral host? Lack of Si. Match to SN 2006oa.
2020qkx	SN Ia	0.127	−19.8	0.89	51	Luminous Ia?
2020opy	TDE	0.159	−20.5	0.89	83	Tagged 23 days before YSE's spectrum.
2021nxq	SLSN-I	0.15	−20.6	0.89	27	Faint host, incorrect assoc.
2020qql	SN Ia	0.076	−19.8	0.88	35	Luminous Ia?
2021hrj	SN Ib	0.022	−17.4	0.88	79	Tagged 7 days before YSE's spectrum (Figure 8).
2021bmv	SN IIn	0.09	−19.3	0.85	11	Faint host.
2021aadc	SLSN-II	0.1953	< -20.8	0.85	15	Faint host.
2020qmj	SN IIn	0.022	−18.8	0.82	51	…
2020kre	SN Ia-CSM	0.136	−19.9	0.82	70	Tagged anomalous 37 days before YSE's spectrum.
2021too	SN Ic-BL	0.07	−19.4	0.82	36	…
2021pnp	SN IIb	0.03	−17.0	0.79	46	…
2021gno	SN Ib-pec	0.0062	−15.0	0.78	29	See W. Jacobson-Galán et al. (2022).
2020kbl	SN Ia	0.079	−18.8	0.77	22	SALT3 $c = +0.3$. In flocculent spiral/merger galaxy.
2021btn	SN II	0.083	−19.6	0.75	43	Bright. No g- obs for ∼50 days. $\rm H_\alpha$ origin unclear.
2021udc	SN IIb	0.035	−18.3	0.75	8	…
2021bpq	SN Ia	0.1	−19.3	0.75	5	…
2021ojn	SN Ia	0.08	−19.5	0.74	18	Faint host, misassoc. Early Ia-bump candidate.
2020tan	SN IIn	0.079	−19.0	0.73	60	Tagged 75 days before YSE's spectrum.
2020acun	SN II	0.0216	−17.3	0.73	25	Misclassified by ParSNIP, SuperRAENN (Ib/c).
2020kvl	SN Ia-91T-like[b]	0.12	−20.6	0.72	8	Possible SN Ia-SC.
2020wfg	SN Ia	0.108	−19.4	0.72	3	…
2020ivg	SN IIb	0.053	−17.3	0.72	40	Tagged 18 days before YSE's spectrum.
2020kpz	SN II	0.039	< -18.2	0.71	65	Preexplosion data? Peculiar light curve (LC). Visible ∼400 days.
2020acct	SN Ic	0.035	−18.0	0.70	32	See Angus et al. (in prep).
2019wmr	SN II	0.038	−17.5	0.70	2	Misclassified by ParSNIP, SuperRAENN (Ib/c).
2021vwx	SN Ia	0.06	< -19.3	0.70	9	No rise or peak. In spiral.
Spectroscopic (Host-z)						
2020rss	SN Ia?	0.1288	−19.1	0.77	11	…
2021ofr	SN Ia?	0.0848	−19.2	0.76	8	…
2020kmj	SN II?	0.0799	−18.4	0.76	49	Elliptical host galaxy.
Photometric (Photo-z)						
2020hjv	SN Ia?	0.141	−19.9	0.91	7	Luminous Ia?
2020acyu	SN IIn?	0.254	−20.4	0.90	33	Faint host.
2020jvi	SN IIn?	0.192	−20.2	0.84	9	Faint host.
2019vuz	TDE?	0.21	−20.8	0.79	17	Blue. Nuclear. FLEET = 89% TDE.
2021dpa	SLSN?	0.152	−18.9	0.73	20	Faint host, misassoc. Peculiar long rise (∼100 days).
2020itp	SN Ia-CSM?	0.16	−19.7	0.73	44	ZTF-r declines slowly but ZTF-g fades quickly.
2021rmq	SN II?	0.131	−19.1	0.73	10	Candidate member of long-rising SN II class.
2020wwt	SN Ia?	0.157	−19.2	0.72	14	SALT3 $c = -0.21$. Spiral. Artifact in PS1-r image.
2020iga	SN Ia?	0.13	−19.2	0.71	7	SALT3 $c = -0.20$. Nuclear.
2020sgy	SN Ia?	0.204	−20.6	0.71	18	Luminous Ia? Faint host.
2020vpn	SN Ia?	0.136	−19.4	0.70	5	…

Notes. The bold text designates a transient event that is likely anomalous, and we consider successfully tagged.

[a] We reclassify from SN Ia to SN Ia-91T-like.

[b] We reclassify from SN Ia to SN Ia-91T-like.

host or small elliptical), two have no rise information (SN 2023dgp as stated before, SN 2023cyx), and one has only r-band detections with g-band nondetections for the last 36 days of the light curve (SN 2023bfv; of which it was tagged anomalous on the last observation overall, meaning that the g-band light-curve features are outdated by several epochs).

Of the photometric anomalies, our model tags 15 objects, of which we consider nine to be likely anomalous. The strongest anomaly candidates are two missed likely TDEs (AT 2022zyh, AT 2023adr), notable for their strong blue color throughout the light-curve duration and their location at the host galaxy nucleus (van Velzen et al. 2019), and our discovery of the long-rising SN II candidate AT 2023inr, which had a rising light curve of at least ∼75 days.

By leveraging `LAISS` and additional tests such as the `iso_forest_anomaly_detection` *Filter*, we can uncover a variety of anomalous objects in the night sky, some of which may have been previously overlooked. This underscores the significance of employing multiple anomaly detection techniques to maximize the scientific value we can extract from recovering known and likely anomalies, quantify common pitfalls for why objects were either missed or not targeted for follow-up, and limit missed opportunities for active study in the future.

4. Retroactive Anomaly Detection with YSE DR1

The premiere multiband, multisurvey time-domain data set is the YSE DR1 (Aleo et al. 2023), which includes the final photometry of 1975 transients observed by the ZTF (Bellm et al. 2019) in *gr* and PS1 (Chambers et al. 2016) in *griz* bands as conducted by the YSE (Jones et al. 2021). YSE DR1 stands out as the most extensive and consistent multiband data set of SNe at low redshifts ever assembled. This data set serves as an ideal testing ground: it features a diverse range of real objects (SN Ia, SN Ia-SC, SN Ia-CSM, SN Ia-91T-like, SN Ia-91bg-like, SN II, SN IIn, SN IIb, SN Ib, SN Ib-pec, SN Ibn, SN Ic, SN Ic-BL, TDE, SLSN-II, SLSN-I, SNa Iax, and rare SN impostors such as a luminous blue variable outburst and a luminous red nova), SNe observations spanning extensive time frames (2019 November 24 to 2021 December 20), and a vast range of apparent magnitudes ($m \in [12, 22]$) and absolute magnitudes ($M_{\rm abs} \in [-13.5, -22.5]$), and covers a redshift distribution up to approximately $z \approx 0.5$. Moreover, because ZTF and PS1 share *gr* photometry, we treat PS1-*gr* and ZTF-*gr* as equivalent by stacking their light curves, effecting reducing the overall cadence and adding depth information.[53]

As an extension of the original work from Aleo et al. (2023), we run `LAISS` to find the most anomalous objects in YSE DR1. We do not use the redshift or latent embedding information from the ParSNIP photometric classifier in our model. Of the 1975 total YSE DR1 data set, 1153 (≈58%) are successfully processed through our model after selection cuts for the requisite number of observations in the combined PS1+ZTF-*gr* bands and associated hosts with PS1-*grizy* photometry. Unlike Tables 4 and 6, we consider the max anomaly score max(P(anom)) at *any* light-curve phase. Note that we use the same RFC model that is trained on the full light curves, but the input YSE DR1 light-curve features are constantly recalculated (the host features do not change) with the `light-curve` extractor and passed through model at each epoch, to mimic a real-time anomaly detection scenario.

There are 198/1153 (≈17%) transients that achieve max(P(anom)) ⩾ 50%, but only 42/1153 (≈3.6%) that achieve max(P(anom)) ⩾ 70%. For brevity, we show the results of these 42 objects in Table 7, ordered by highest anomaly score. If two or more objects have the same score, we secondarily rank them by the number of observations for which the anomaly score of the object was P(anom) ⩾ 50% ("Anom. Obs").

When considering *spectroscopic* anomalies alone, we find a purity of 15/29, or 52%, aligning well with the results we find in Figure 5. If expanded to consider *behavioral* and *contextual* anomalies, we achieve a purity of 20/29, or 69%, which slightly underperforms given what we expect from Figure 5. However, our anomaly detection model is not trained nor tested with PS1 photometry nor ZTF PSF-fit forced-photometry[54] (as opposed to the ZTF Alert Stream) as in YSE DR1, and thus, we cannot expect our model to perform in exactly the same way.

The original work from Aleo et al. (2023) did not do an extensive study of anomalous objects. Broadly, they found transients with latent embeddings (from the YSE-ZTF trained ParSNIP classifier) outside the bulk of their member class distribution were more likely to be misclassified. They indicated that these objects likely deviated from the normal, more representative objects that share their spectroscopic class. Because the ParSNIP classifier was trained on simulations from idealized templates, objects whose real behavior differed from this normality were likely to be poorly characterized. Despite classifying SNe Ia with high purity (⩾90%), they found misclassified SNe Ia tended to fall within two categories: (1) observing effects such as significant (∼100 day) gaps, only observed well after peak (e.g., SN 2020zmi, SN 2021van, SN 2021vwx), and SN Ia that requires a large correction for dust extinction (e.g., red SNe Ia with SALT3 $c > 0.3$; SN 2020pki, SN 2020zfn, SN 2021aamo); or (2) physical effects linked to rare phenomena (91T-like/91bg-like) properties (SN 2021bmu, SN 2021ctn), and long-lived CSM interaction (e.g., SN Ia-CSM: SN 2020aekp, SN 2020kre, SN 2021uiq).

If we investigate misclassified SNe Ia due to observing effects, we find the following: the hosts of SN 2020pki and SN 2020zmi did not have the requisite *grizy* photometry, and thus were not processed by our model; SN 2021van did not have enough *g*-band photometry; SN 2020zfn and SN 2021aamo were successfully processed by our model, but not tagged anomalous; SN 2021vwx was tagged anomalous, although likely due to having no rise nor peak information in the light curve, which is rare in our training set (see Section 6.2 for details). We have insufficient statistical evidence to determine how YSE DR1 SNe affected by observing effects influence our `LAISS` anomaly detection classification.

We additionally investigate objects misclassified by ParSNIP and SuperRAENN, as these are strong anomaly candidates. Note that both classifiers in Aleo et al. (2023) are only trained via simulations to classify into SN Ia, SN II, and SN Ib/c (thus, SLSN, TDE, and more will always be misclassified). We find 13 such objects flagged by our anomaly detection model:

1. *SN 2020xsy (SLSN-II).* This is the most anomalous object from YSE DR1 according to our model, in terms of max(P(anom)) at 98% and number of epochs with an anomaly classification (93), and also, at the furthest redshift, $z = 0.27$. ParSNIP misclassified as SN II; SuperRAENN misclassified as SN Ia.
2. *TDE 2020opy (TDE).* This was tagged anomalous 23 days before the first classification spectrum from YSE, and the second highest number of epochs with an

---

[53] While this approximation is generally accurate, minor discrepancies may arise due to variations in passband transmission profiles and disparities in photometric pipelines. Moreover, ZTF is not color calibrated; it is calibrated to PS1 assuming that all ($g - r$) colors are zero. We do not account for the ZTF color correction coefficient. Therefore, the more rigorous approach involves considering passbands independently, which is left for future work. See http://nesssi.cacr.caltech.edu/ZTF/Web/gettingto1.html for more details.

[54] https://web.ipac.caltech.edu/staff/fmasci/ztf/forcedphot.pdf

anomaly classification (83). ParSNIP misclassified as SN Ia; SuperRAENN misclassified as SN Ibc.

3. *SN 2021nxq (SLSN-I).* Even with an incorrect host association (likely due to the faint host), this object is tagged anomalous at peak light. ParSNIP misclassified as SN Ib/c; SuperRAENN misclassified as SN Ib/c.
4. *SN 2021aadc (SLSN-II).* This was tagged anomalous 2 days before the first classification spectrum from BTDG, approximately 2 weeks before peak light. There is no decline information for this object. ParSNIP misclassified as SN Ib/c; SuperRAENN misclassified as SN Ib/c. It is located in a faint host.
5. *SN 2020qmj (SN IIn).* From peak light onwards, there is an approximately constant red color $g - r \approx 0.7$ mag in the light-curve evolution. ParSNIP misclassified as SN Ia; SuperRAENN misclassified as SN Ia. It is located in an edge-on spiral galaxy.
6. *SN 2020kre (SN Ia-CSM).* Our anomaly detection algorithm tagged this object as anomalous 21 days after the original SN Ia-normal classification but 37 days *before* YSE's SN Ia-CSM reclassification spectrum. It is located in a faint, blue host.
7. *SN 2021pnp (SN IIb).* From peak light onwards, there is an approximately constant red color $g - r \approx 0.9$ mag in the light-curve evolution, photometrically similar to SN 2020qmj. There is no strong evidence of shock breakout in the photometry. ParSNIP misclassified as SN Ib/c; SuperRAENN misclassified as SN Ib/c.
8. *SN 2021btn (SN II).* This is unusually bright for a normal SN II at $M_{\rm abs} \sim -19.6$ mag. There were a lack of g-band observations for the first 50 days of the event. It is unclear if the H_α emission in the spectrum is from the host or the transient, and whether there are signs of P-cygni. A possible alternative is an SN IIn classification. ParSNIP misclassified as SN Ib/c; SuperRAENN misclassified as SN Ib/c.
9. *SN 2021udc (SN IIb).* There is no strong evidence of shock breakout in the photometry. ParSNIP misclassified as SN Ia; SuperRAENN misclassified as SN Ia.
10. *SN 2020acun (SN II).* From peak light onwards, there is an approximately constant red color $g - r \approx 0.9$ mag in the light-curve evolution, photometrically similar to SN 2020qmj and SN 2021pnp. `Superfit` (Howell et al. 2005) gives matches to SN IIb, but the He I 5875 Å absorption is very low signal-to-noise ratio if it is present at all, and hence, we keep the SN II classification. ParSNIP misclassified as SN Ib/c; SuperRAENN misclassified as SN Ib/c.
11. *SN 2020kpz (SN II).* This is a peculiar SN with possible preexplosion data, marked by a sudden shift from blue to red before the sharp rise to peak. ParSNIP misclassified as SN Ib/c; SuperRAENN misclassified as SN Ib/c.
12. *SN 2020acct (SN Ic*).* This is a peculiar, distinctly double peaked SN separated by ~60 days, which does not fall neatly into any existing SN classification. See Angus et al. in prep for details. ParSNIP misclassified as SN Ia; SuperRAENN misclassified as SN Ib/c.
13. *SN 2019wmr (SN II).* From peak light onwards, there is an approximately constant red color $g - r \approx 0.9$ mag in the light-curve evolution, photometrically similar to SN 2020qmj, SN 2021pnp, and SN 2020acun. ParSNIP misclassified as SN Ib/c; SuperRAENN misclassified as SN Ib/c.

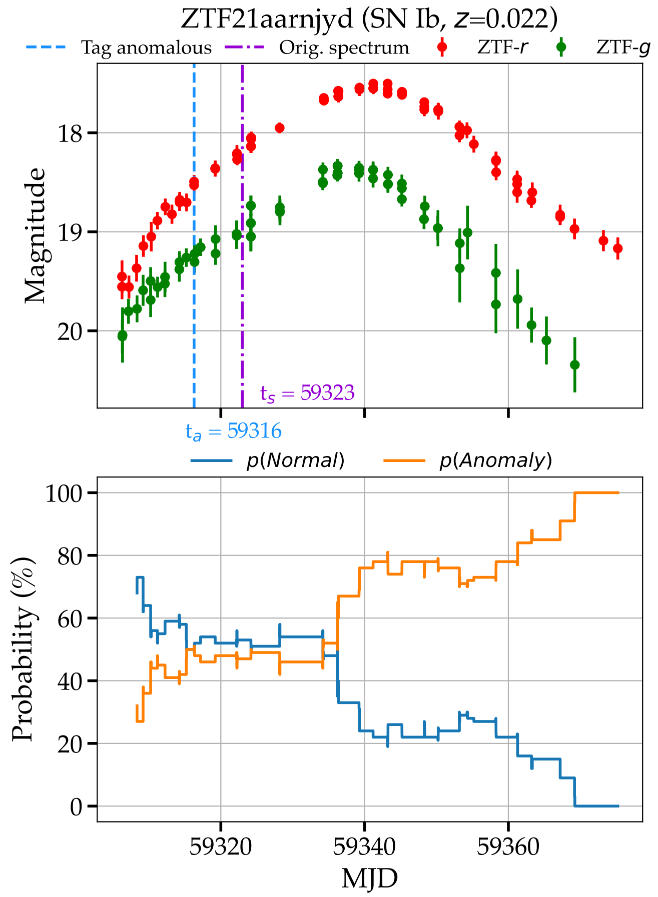

Figure 8. The Type Ib SN 2021hrj/ZTF21aarnjyd was retrospectively tagged as anomalous (P(anom) ⩾ 50%) on MJD 59316 (2023 April 12) by our `LAISS` AD algorithm (dashed blue line), *1 week earlier* than the first classification spectra from YSE (2023 April 19; dotted–dashed purple line) and ZTF (2023 April 20) on TNS. The YSE team originally requested spectra when this object passed their magnitude-limited sample criteria. If our algorithm was in place when SN 2021hrj was active, we could in principle have obtained a spectrum 1 week earlier in the light-curve evolution, providing valuable information in regards to the early time supernova physics and local environment of this peculiar, long-rising event.

An interesting observation is that, of the 28 most anomalous spectroscopic YSE DR1 objects (Table 7), 13 (or 46%) are what we consider bona fide anomalies (bold text) *and* are misclassified by both ParSNIP and SuperRAENN. When accounting for only spectroscopic classes for which ParSNIP and SuperRAENN were trained (SN Ia, SN II, SN Ib/c; and thus excluding SLSN and TDE), we find this drops to only nine (or 32%).

Although we do not formally consider the phase at which the anomaly detection successfully tags an object as anomalous in our metrics, there is great value at flagging anomalies at or before peak, and as early time as possible, as these times are when the physics of the progenitor explosion is most correlated with the photometry (Gagliano et al. 2022). Of the successfully tagged objects in Table 7, our algorithm correctly identifies 6/20 (30%) as anomalous *at least* 1 week earlier than the first classification spectrum. In the best case, we identified SN IIn 2020tan 75 days prior. We show the "worst" case of only 7 days prior for SN Ib 2021hrj in Figure 8. The YSE team originally requested spectra when SN 2021hrj passed their magnitude-limited sample criteria. If our anomaly detection algorithm were in place when it was active, we could, in principle, have obtained a spectrum 1 week earlier in the light-curve evolution, providing valuable information in regards to earlier-time SN physics and local environment of this peculiar, long-rising event.

We additionally identify strong anomaly candidates for which there was no spectroscopic follow-up (Table 7). We consider such photometric objects as missed opportunities. Although it is impossible to know the exact classifications, we have evidence that 6 of 11 (55%) of the most anomalous tagged photometric objects from YSE DR1 were likely bona fide anomalies:

1. *AT 2020acyu (SN IIn?).* This is visible for ∼175 days, and approximately 2.″42 separation from the center of faint host PS1 galaxy catalog ID 1237654653639656020. At an estimated redshift of $z = 0.254$ from `Easy Photoz` (Aleo et al. 2023), we calculate a peak absolute magnitude of $M_{\mathrm{abs}} \approx -20.4$ mag. Likewise, if we use the photo-z estimate of $z = 0.224$ from SDSS (Csabai et al. 2003), we estimate a peak absolute magnitude of $M_{\mathrm{abs}} \approx -20.2$ mag. FLEET (Gomez et al. 2020) predicts a Type II (with no distinction between Type II and Type IIn) with 68% confidence, followed by SLSN-II at 24%. ParSNIP predicts a Type II with 55.9%. With the luminous but not superluminous brightness, extended duration, and color evolution, we predict this to be a missed SN IIn. Beyond the simple explanation of being overlooked, a possible reason for why there was no spectroscopic follow-up requested for this object is because this object achieved a peak apparent magnitude of only ∼20.0 mag, which would require long exposure times for such a faint source, and thus would be a low priority.
2. *AT 2020jvi (SN IIn?).* This is visible for ∼120 days, approximately 0.″11 separation from the center of faint, blue host PS1 galaxy catalog ID 1237662636371739674. At an estimated redshift of $z = 0.192$ from `Easy Photoz` (Aleo et al. 2023), we derive a peak absolute magnitude of $M_{\mathrm{abs}} \approx -20.2$ mag. Likewise, if we use the photo-z estimate of $z = 0.173$ from SDSS (Csabai et al. 2003), we calculate a peak absolute magnitude of $M_{\mathrm{abs}} \approx -20.0$ mag. This object rose approximately 2 mag in 20 days, and is marginally red (∼0.1 mag) throughout its duration. FLEET (Gomez et al. 2020) predicts an SLSN-II at 58% confidence. To add to the confusion, ParSNIP predicts a Type Ib/c at 69% confidence (however, ParSNIP is not trained for SLSN classification). An SALT3 fit reveals a poor fit and a stretch parameter value of $x_1 = +3.00$. Similarly as AT 2020acyu, with the luminous but not superluminous brightness, extended duration, and color evolution, we cautiously predict this to be a missed SN IIn. This object achieved a peak apparent magnitude of ∼19.4 mag in ZTF-r, which is within reasonable limits for follow-up.
3. *AT 2019vuz (TDE?).* Arguably the most anomalous of the photometric subset due to the intrinsic rarity, we predict this object as a likely missed TDE, due to its duration of ∼75 days, blue color, at a nuclear location of 0.″11 separation from the center of its red, possible elliptical host PSO J072914.403+420437.05. At an estimated redshift of $z = 0.21$ from `Easy Photoz` (Aleo et al. 2023), we estimate a peak absolute magnitude of $M_{\mathrm{abs}} \approx -20.8$ mag. Using the estimated redshift of $z = 0.248$ from PS1-STRM (Beck et al. 2021), we estimate a peak absolute magnitude of $M_{\mathrm{abs}} \approx -21.2$ mag. FLEET (Gomez et al. 2023d) predicts a TDE at 89% confidence. This object achieved a peak apparent magnitude of ∼19.3 mag in PS1-g and was brighter than 19.5 mag for ∼20 days, which is well within reasonable limits for follow-up.
4. *AT 2021dpa (SLSN?).* Exhibits an unusually prolonged slow rise of ∼0.8 mag over ∼100 days (estimated ∼125 days mag^{-1}) in an extremely faint, small, red host. Likely due to its faint nature, `GHOST` (Gagliano et al. 2021) does not find the likely host and results in a misassociation (to a nearby, faint galaxy PS1 object ID 1237664871896777642). At approximately 2.″2 offset from its apparent host PS1 object ID 149701403944978764, there are two disparate estimates for the photo-z. From `Easy Photoz` (Aleo et al. 2023), the estimated redshift of $z = 0.152 \pm 0.096$ results in a peak absolute magnitude of $M_{\mathrm{abs}} \approx -18.9$ mag. However, the YSE DR1 estimate is likely low, due to the host being perceptible in the DESI Legacy Survey Imaging Surveys[55] but only marginally so in SDSS and PS1 stacked templates. The SDSS estimated redshift of $z = 0.482 \pm 0.1461$ (Csabai et al. 2003) results in a peak absolute magnitude of $M_{\mathrm{abs}} \approx -21.9$ mag. FLEET (Gomez et al. 2023a) estimates an SLSN-II with 40% confidence, followed by SLSN-I with 31% confidence. When considering proprietary YSE photometry beyond the YSE DR1 cutoff, this object is visible for ∼370 days. All of this evidence together points to an SLSN. A compelling reason for why there was no spectroscopic follow-up requested for this object is because this object achieved a peak apparent magnitude of only ∼20.2 mag before disappearing behind the Sun, only to reappear at a faint 21.7 mag at ∼250 days later. It is unlikely one could acquire spectroscopic follow-up for this object even if it was known to be anomalous.
5. *AT 2020itp (SN Ia-CSM?).* Visible for ∼80 days, in elliptical host PS1 objectID 147052397040217754. The ZTF-r photometry declines slowly but the ZTF-g fades quickly. There is a fair agreement in photo-z. From `Easy Photoz` (Aleo et al. 2023), the estimated redshift of $z = 0.16 \pm 0.051$ results in a peak absolute magnitude of $M_{\mathrm{abs}} \approx -19.7$ mag whereas the SDSS estimated redshift of $z = 0.109 \pm 0.033$ (Csabai et al. 2003) results in a peak absolute magnitude of $M_{\mathrm{abs}} \approx -19.0$ mag; both are consistent with SN Ia. Regardless, for both photo-z estimates, if there existed any CSM interaction, the H_{α} profile would reside in only the ZTF-r wavelength range,[56] a possible driver for the slow ZTF-r decline. Because this object is as intrinsically bright as an SN Ia but with a longer duration and slower ZTF-r decline that is consistent with possible CSM interaction, we posit this to be a missed SN I-CSM. This object achieved a peak apparent magnitude of ∼19.4 mag in ZTF-r, which is within reasonable limits for follow-up.
6. *AT 2021rmq (long-rising SN II).* The YSE DR1 light curve shows a rise of ∼60 days, with an evolution consistent with Type II, and is thus a candidate member

---

[55] https://www.legacysurvey.org/

[56] http://svo2.cab.inta-csic.es/svo/theory/fps3/index.php?id=Palomar/ZTF.r&&mode=browse&gname=Palomar&gname2=ZTF#filter

> of the rare long-rising SN II class (Sit et al. 2022). Additional ZTF DR data photometry[57] indicate that this rise may be longer, closer to ∼80 days. It is offset 0.″23 from its blue (likely spiral) host PS1 objectID 146701848390420314. The `Easy Photoz` (Aleo et al. 2023) photo-z (0.131) is in good agreement with that of Beck et al. (2021; 0.146), placing the peak absolute magnitude at around $M_{\rm abs} \approx -19.1$ mag. However, if the photo-z is accurate, this would place the absolute magnitude brighter than any of the 13 events found in Sit et al. (2022) at −17.5 mag. It is also possible the long rise could be a result of CSM interaction (Nyholm et al. 2020), and this object could instead be an SN IIn.

The remaining five photometric objects with high anomaly scores (AT 2020hjv, AT 2020wwt, AT 2020iga, AT 2020sgy, AT 2020vpn) are likely SN Ia. If photo-z estimates are correct, two are luminous and potentially of a rare subtype at $M_{\rm peak} \sim -20$ mag (AT 2020hjv, AT 2020sgy), but we find no conclusive evidence. Of the others, two are blue when compared to the SALT3 (Kenworthy et al. 2021) model at $c \sim -0.2$ (AT 2020wwt, AT 2020iga), and one (AT 2020vpn) appears to be unremarkable.

Lastly, there are three objects with host spectra but no transient spectra. AT 2020rss and AT 2021ofr appear to be normal SN Ia based on evolution and absolute magnitude estimates using the spectroscopic host-z. The likely *contextual* anomaly is AT 2020kmj, which exhibits an SN IIP profile with typical SN II peak absolute magnitude but resides in what appears to be an elliptical galaxy (PS1 objectID 116361881748977931).

Overall, of the objects that were successfully processed and characterized by `LAISS` from YSE DR1 (1153/1975, or 58%), 7/27 (26%) of the most anomalous objects (P(anom) ⩾ 70% at any phase in the light curve) that we consider likely anomalies (bold text) did *not* have transient spectra. Of those that did, 6/20 (30%) were flagged by `LAISS` at least a week prior to the classification spectrum.

5. Approximate Nearest Neighbor Similarity Search with ANNOY

In this current age of large time-domain surveys like ZTF (Bellm et al. 2019) and the imminent LSST (Ivezić et al. 2019), the bottleneck of object discovery lies in the ability of automated algorithms to quickly isolate and prioritize detections of interest amid the flood of alert events. ZTF generates approximately 70 GB night^{-1} (Patterson et al. 2019) and discovers an order ∼10,000 SNe yr^{-1},[58] and LSST is estimated to produce an alerts volume of 10× that at approximately 782 GB night^{-1}, observing an estimated order ∼1 million SNe yr^{-1}.[59]

With high present-day discovery rates and climbing, we cannot rely on computationally expensive operations for SNe discovery and characterization. In this section, we argue that a low-latency, sublinear time-complexity ANNs search can be an effective tool to search for similar SNe for applications of SN discovery, reclassfication, and more.

Previous studies have shown that data mining strategies involving a nearest neighbors search for retroactive SN discovery can be successful (Aleo et al. 2022). However, Aleo et al. (2022) used a brute-force kD-tree approach to search for neighbors in feature-space, which in the worst case scales linearly with the number of data points in runtime ($\mathcal{O}(n)$). This approach would be impractical for the real-time data streams of ZTF and LSST. We use their approach as inspiration, and instead opt for an approximate similarity search, which in the worst case scales in logarithmic time with respect to the number of data points, $\mathcal{O}(\log(n))$.

For our ANN algorithm, we use the open-source package `ANNOY`[60] (Bernhardsson 2018), the method developed by Spotify for song recommendations. Essentially, `ANNOY` allows for an efficient ANN search in high-dimensional spaces, such as the large light curve and host galaxy feature data set we use in this work.[61] This process involves dividing the high-dimensional space into smaller subspaces using random hyperplanes. These subspaces are organized into a binary tree structure, where each node represents a subspace, and each leaf node corresponds to a specific point in the original high-dimensional space. During a search operation, the algorithm begins at the root of the tree and recursively traverses the tree, selecting the branch that is closest to the query point. At each leaf node, the distance between the query point and the represented point is recorded. Simultaneously, a list of the k closest points encountered thus far is maintained. To optimize the search process, recorded distances are utilized to prune branches unlikely to contain points closer than the current kth closest point. This iterative search is repeated for each query point. The `ANNOY` library is designed to ensure efficient memory usage and high performance, enabling indexing of millions of high-dimensional vectors.

In the `LAISS` pipeline (Figure 1), we build an `ANNOY` index of our databank—the reference data set to which new, incoming, or manually chosen objects will be compared. However, we do so after scaling our databank with the `sklearn.preprocessing.StandardScaler()`[62] module and applying PCA (Jolliffe 2002) to reduce the dimensionality of our data from 120 dimensions to 60 dimensions, retaining ∼98% of the variance. This will partially speed up the ANN search without important information loss; moreover, having less than 100 dimensions is recommended by the code authors for the best performance.[63]

`LAISS` can run the ANN search to calculate any k neighbors (default is $k = 8$) for any ZTF object for which its light-curve and host galaxy association and features can be extracted, regardless if it exists in the databank or not; all that is needed is a ZTF object ID of the user's choosing. For example, when supplied a ZTF object ID, `LAISS` queries ANTARES for the object of interest's ZTF photometry and manually extracts the light-curve features in the same manner as the `lc_feature_extractor` *Filter*, followed by running `GHOST`

---

[57] https://www.ztf.caltech.edu/ztf-public-releases.html

[58] For this estimate, we use the search function from www.wis-tns.org/ and search for all transients reported to TNS from ZTF first light on 2017 November 1 to 2023 November 1 with ZTF listed as the discovery source, and find the per year average.

[59] https://dmtn-102.lsst.io/DMTN-102.pdf

[60] https://github.com/spotify/annoy

[61] Note that in this work we do not use data imputation. Although in principle `ANNOY` can run on data sets with imputed entries, it may introduce bias and impact the nearest neighbor search accuracy, particularly if many values are imputed.

[62] https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.StandardScaler.html

[63] Though, in principle, it can perform "surprisingly well" up to 1000 dimensions; see https://github.com/spotify/annoy#summary-of-features.

(Gagliano et al. 2021) to retrieve the matched-host galaxy and its features. The object will undergo the same scaling and PCA transformation as the data set bank, and then `ANNOY` will find the ANNs in the 60-dimensional PC-space. As the ANNs are found, `LAISS` queries the ANTARES broker to retrieve the up-to-date ZTF photometry, IAU name, spectroscopic classification (if exists), and spectroscopic redshift (if exists) from TNS for each ANN. Plotting functions are then run to overlay the reference transient and its k ANNs' light curves and host galaxy thumbnails. Currently, this ANN similarity search feature is only available in the Python module on Github.[64] Efforts to create the same functionality within an ANTARES *Filter* are ongoing. This will remove the additional outside API requests for extracting light-curve and host galaxy features and up-to-date TNS information.

Within our databank of 5472 transients, the default $k = 8$ ANN search alone takes anywhere from 1 to 350 ms on a 2 GH. Quad-Core Intel Core i5. This results in $\leqslant$5 minutes of total processing time to find $k = 8$ ANNs for all $\sim$1000 SN night^{-1} estimated with LSST. If an object needs to have its light-curve and host features extracted from scratch, the entire process takes about 1 minute per transient (thus, $\sim$16 hr of total processing time for all $\sim$1000 SN night^{-1}), where the time bottleneck stems from the `GHOST` host association—specifically time spend on catalog queries (e.g., PS1); such catalogs could be downloaded to disk to dramatically speed up associations. Utilizing faster host association methods such as DELIGHT ($\sim$60 ms/transient, Förster et al. 2022) and modifying `GHOST` to use internal disk-downloaded catalogs is a subject of future work.

5.1. ANN Similarity Search Results on Light Curves and Host Galaxies

The ability to quickly find analogs to an object of interest is crucial to many useful applications, including but not limited to the following:

1. *Predicting behavior*. Based on the known properties and evolution of analogs, one can predict the potential behavior of newly observed transients across their time evolution. Nominal events will likely evolve as expected in common galactic environments, and deviations from this behavior (or in the case where analogs are anomalous) could indicate rare or misunderstood phenomena;
2. *Calibrating models*. By comparing observations of known analogs with theoretical predictions, experts can improve the accuracy of models describing the underlying physical processes of transient events, such as those used for simulations (e.g., models developed for `SNANA`, Kessler et al. 2009; and their use in PLAsTiCC, Hložek et al. 2023).
3. *Studying rare events*. By associating a similarity score, experts can measure how similar or dissimilar an event is from the rest of the existing data set. Objects strongly dissimilar to their nearest neighbors can indicate truly unusual phenomena, or out-of-distribution (OOD) behavior for the particular reference data set.
4. *Probing host galaxy properties*. Analogous transients in different host galaxies allow experts to investigate the influence of host galaxy properties on the characteristics of the transient events. This information is valuable for understanding the connections between stellar populations, galactic environments, and transient phenomena.
5. *Improving classification*. As more analogs are discovered and characterized, classification systems can be refined, leading to more accurate and automated identification of new transients.
6. *Reclassification*. A given transient of a known type whose analogs are classified differently can direct the expert's attention into reinspecting their spectra and classifications, possibly updating their classification. We show this to be true in Section 5.2 and Table 8, where we update classifications for 17 unique SNe from our databank, prompted by an ANN search.
7. *Optimizing observational strategies*. Observers can prioritize targets based on the likelihood of finding analogs or because they are analogs to a known rare or scientifically useful event, helping to maximize the efficiency of observing campaigns and the utilization of telescope time. This can be used in tandem with active learning strategies like RESSPECT (Kennamer et al. 2020).
8. *Retroactive discovery*. Analogs of known transients can also be used as a data mining method to find similar but previously undiscovered (i.e., not publicly reported) transients in large data sets, successfully demonstrated by Aleo et al. (2022). In this work, we showcase 84 transients discovered from the ZTF Alert Stream in 2018–2021 using this method, as shown in Table F1 and Section 5.4. Similarly, analogs can be used to suggest likely candidates of a certain class that were originally reported but have no follow-up spectra for classification, and thus represent possible missed opportunities (Section 5.3).

To our knowledge, there is no popular and lightweight tool for finding transient analogs like SNe. To demonstrate such a utility, we present an ANN similarity search of the most commonly observed SNe (a Type Ia) in Figure 9. As a reminder, we only input the 60-dimensional PC feature-space (from the 120-dimensional light curve and associated host galaxy feature-space). We *do not* input the photometric observations themselves, nor the redshift, SN type, galaxy type (spiral, elliptical), or thumbnail image (although this is a topic to be explored in future work). For visual aid, in the left panel, we overlay the object of interest (ZTF21aaublej/SN 2021ixf) and its $k = 8$ closest ANNs, shifting the light curves such that the peak brightness in the ZTF-r band and ZTF-g band is at 0 day since peak (both passbands fit independently). In the right panel, we show the thumbnail of the host galaxy of our object of interest in the top left, and the eight ANNs continuing from left to right, top to bottom. This similarity search in our database was completed in 3 ms.

We find that the reference SN 2021ixf, in part because it is a well-sampled, full-phase, common SN Ia at a nearby redshift, is indeed highly similar in light-curve evolution and host galaxy environment as its eight ANNs. All light curves increase in brightness by $\approx$3 mag in 16–20 days to peak from their first observation, and have a prominent secondary ZTF-r-band bump at around 30 days postpeak, lasting about 100 days in

---

[64] https://github.com/patrickaleo/LAISS-local

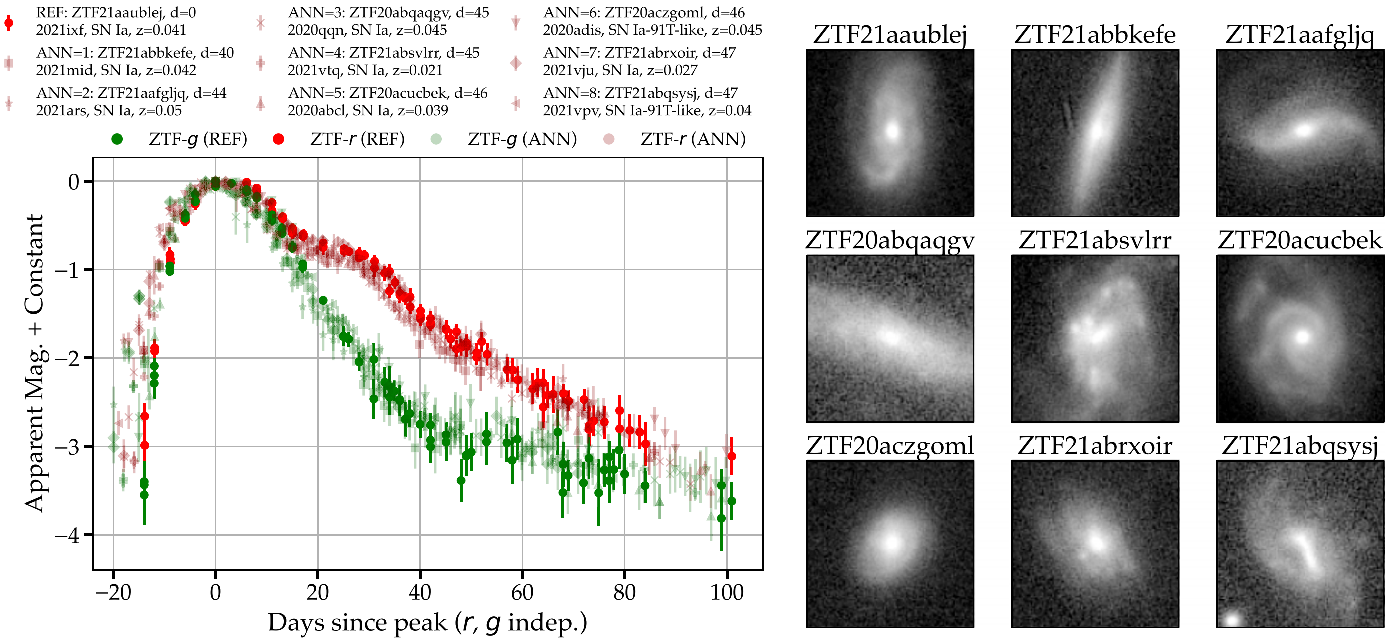

Figure 9. Stacked light curve (left panel) and grid of host galaxy thumbnails (right panel) of Type Ia SN 2021ixf (bold) and its eight approximate nearest neighbors within 120 day feature-space (faded). The similarity search in our database was completed in only 3 ms. Note that we do not include SN type or redshift information in the similarity search. Thus, the similar light-curve evolution, SN types, redshift values, and host galaxy environments are as a result of the close proximity of SN 2021ixf and its neighbors in 120-dimensional feature-space.

Table 8

Reclassification of SNe Updated Classification of 17 Unique SNe Driven by ANN = 8 Nearest Neighbor Matches to All SNe Classes Using `SNID`

ZTF ID	IAU Name	TNS Class	New Class	Reference SNe	ANN	Remarks
Reference Class: SN IIn (59 objects)						
ZTF21abcjpnm	2021njo	SN II	SN IIn	2020abku	1	Second peak likely CSM interaction.
ZTF21aaizyqc	2021ckb	SN II	SLSN (He-rich)	2021hur	2	Match to PTF10hgi (Gal-Yam 2019).
Reference Class: SN Ia-91T-like (37 objects)						
ZTF21aagoliy	2021cjc	SN Ia	SN Ia-91T-like	2020zjv	1	Match to SN 1997br, SN 1991T.
ZTF21abjtqyq	2021sis	SN Ia	SN Ia-91T-like	2020acef	1	Match to SN 1999aa.
ZTF21abothvq	*2021uib*	SN Ia	SN Ia-91T-like	“	6	Matches to SN 2001V, SN 1991T.
ZTF21abicgai	2021sju	SN Ia	SN Ia-91T-like	2020adis	3	Match to 1997br.
ZTF21abcxner	2021nxh	SN Ia	SN Ia-91T-like	2021qvg	8	Lack of Si. Match to SN 1991T.
Reference Class: SN Ib (25 objects)						
ZTF21aaqwfqe	2021hen	SN I	SN Ib	2021gno	2	Match to iPTF13bvn.
ZTF21aabyifm	2021qv	SN Ib/c	SN Ib	”	5	Matches to iPTF13bvn, SN 2009iz.
Reference Class: SN IIb (21 objects)						
ZTF21abnvlnj	2021tyf	SN II	SN IIb	2021M	8	LC has shock-cooling peak.
Reference Class: SN Ic-BL (14 objects)						
ZTF21aacufip	2021vz	SN Ic	SN Ic-BL	2021too	4	Match to SN 2007ce. Bad quality spectrum.
Reference Class: SN Ia-pec (10 objects)						
ZTF20ackkejs	2020xyd	SN Ia	SN Ia-91T-like	2021cky	8	Lack of Si and bright ($M_{\mathrm{abs}} = -19.65$).
ZTF21abothvq	*2021uib*	SN Ia	SN Ia-91T-like	2021ebb	1	Matches to SN 2001V, SN 1991T.
ZTF21abiawpf	2021rce	SN Ia	SN Ia-91T-like	”	5	Matches to SN 2006cz, SN 2007S.
Reference Class: SN Ia-91bg-like (six objects)						
ZTF20acnznol	*2020yje*	SN Ia	SN Ia-91bg-like	2021jvp	7	Matches to SN 2007ba, SN 2000cn.
”	”	SN Ia	SN Ia-91bg-like	2021wzb	6	”
ZTF21achjwus	*2021abpz*	SN Ia	SN Ia-91bg-like	2021jvp	5	Matches to SN 1986G, SN 1999bh.
”	”	SN Ia	SN Ia-91bg-like	2021wzb	8	”
ZTF21acfigoo	2021aazj	SN Ia	SN Ia-91bg-like	2021fnr	2	Matches to SN 1986G, SN 2007ax.
ZTF21abmwgow	2021ttg	SN Ia	SN Ia-91bg-like	2021wzb	2	Matches to SN 2008R, SN 2007ap.

Note. Italicized text designates transients that appear at least twice in the table. The double quotation ” marks a repeated entry from the above row.

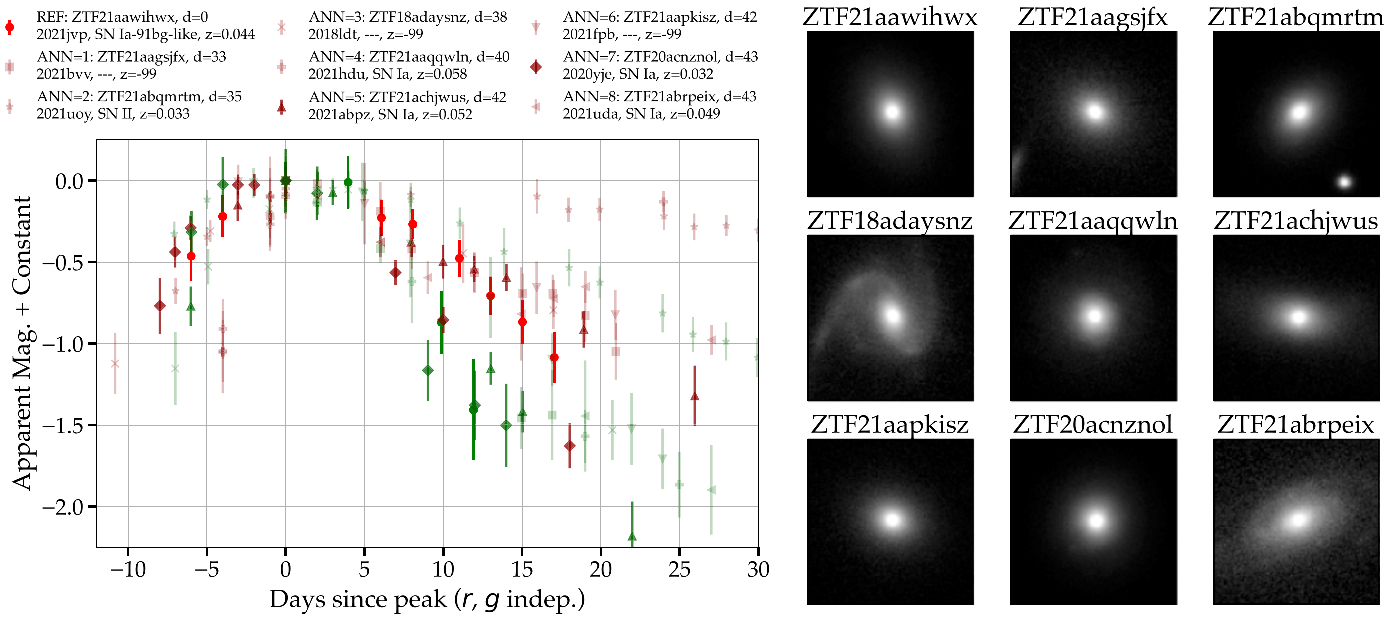

Figure 10. Same as Figure 9, but for Type Ia-91bg-like SN 2021jvp/ZTF21aawihwx. In this case, ANN = 5 (SN 2021abpz/ZTF21achjwus) and ANN = 7 (SN 2020yje/ZTF20acnznol) are in fact SN Ia-91bg-like SNe previously missclassified as SN Ia-normal after reexamination of their TNS classification spectra, prompted by the results of this similarity search (see Table 8). We bold these reclassified events (dark green and dark red for ZTF-g and ZTF-r, respectively) for visual purposes. The similarity search in our database was completed in only ∼300 ms.

total. Our reference object of class SN Ia-normal at $z = 0.041$ is matched to six SN Ia-normal and two SN Ia-91T-like spanning $z \in [0.021, 0.05]$, with a mean redshift of $\bar{z} = 0.039$. The host galaxies are all nearby, face-on spiral galaxies except for one, which is edge-on (SN 2021mid/ZTF21abbkefe). Another also exhibits evidence of a violent merging history (SN 2021vtq/ZTF21absvlrr).

5.2. Reclassification of SNe

Here, we briefly test the ansatz that similar transient types broadly have similar light-curve evolutions and host galaxy environments, investigating whether or not it proves fruitful to look for nearest neighbors of rare reference objects. Although there is no guarantee that nearest neighbors of an object are of the same type, it is a higher likelihood than random sampling. In this search, we find that in doing so we additionally recover transients of rare types that were initially misclassified.

In Figure 10, we show the Type Ia-91bg-like SN 2021jvp/ZTF21aawihwx and its eight ANNs, achieved in ∼300 ms. Initially, if we look at the original TNS classification labels only, we find four SN Ia, one SN II, and three with no classification spectrum (for which we make no definitive determination). Prompted by the nearest neighbor results, we manually reinspect the classification spectra from TNS for these five objects using the Supernova Identification package (`SNID`; Blondin & Tonry 2007). We use the 5.0 version of `SNID` with additional template sets from the Berkeley Supernova Ia Program (Silverman et al. 2012; Liu & Modjaz 2014; Modjaz et al. 2014, 2016; Liu et al. 2016; Gutiérrez et al. 2017; Williamson et al. 2019) totaling 6145 spectra from 811 templates. We use the *forcez* argument for any object that has a known host-z.

We find SN 2021abpz (ANN = 5) and SN 2020yje (ANN = 7) are better explained by the SN Ia-91bg-like classification—the same as reference SN 2021jvp. In SN 2021abpz, there is a visible titanium trough at 4200 Å (Hachinger et al. 2009; Heringer et al. 2017) and strong `SNID` matches to subluminous/Type Ia-91bg-like SN 1986G, SN 1999bh a few days after peak, in phase with when the spectrum was taken. SN 2021abpz has a peak absolute magnitude of $M \sim -18.5$ mag, consistent with subluminous Ia. Similarly, SN 2020yje has strong matches to Type Ia-91bg-like SN 2007ba and SN 2000cn a few days before peak, with a peak absolute magnitude of $M \sim -17.6$ mag, also consistent with subluminous Ia. Moreover, both SN 2021abpz and SN 2020yje are in elliptical hosts and offset from the center—common for Type Ia-91bg-like events (Barkhudaryan et al. 2019). Their light curves are red and fast declining. We additionally find both objects are among the first eight ANNs of another SN Ia-91bg-like object, SN 2021wzb.

We repeat this procedure and manually inspect the TNS classification spectra of all classified transients within the first eight ANNs for the spectroscopic objects in our databank. Any updated classifications for classes in which at least one discovery was made is shown in Table 8. In the majority of cases, the original classification labels are correct, and we do not reclassify. Of our entire spectroscopic sample of 1620 objects, we find 17 unique objects that are better explained by a different classification label, or about 1%. Although this percentage is small, of the 17 objects we reclassify, 8 (47%) were reclassified from a nonanomalous class (which in this work are SN Ia, SN Ia-91T-like, SN II, and SN IIP) to an anomalous class. This brings the total anomalous spectroscopic sample from 228 objects to 236, an increase of 3.5%. Because of the intrinsic rarity of anomalous objects, a 3.5% increase in our overall spectroscopically anomalous sample from a directed manual reevaluation of their spectra is a marked difference.

In our original spectroscopic sample of 1620 objects, only six were originally of this SN Ia-91bg-like classification. After searching this same data set for the eight ANNs of the six known SN Ia-91bg-like objects, we find an additional four SN Ia-91bg-like objects originally misclassified as SN Ia-

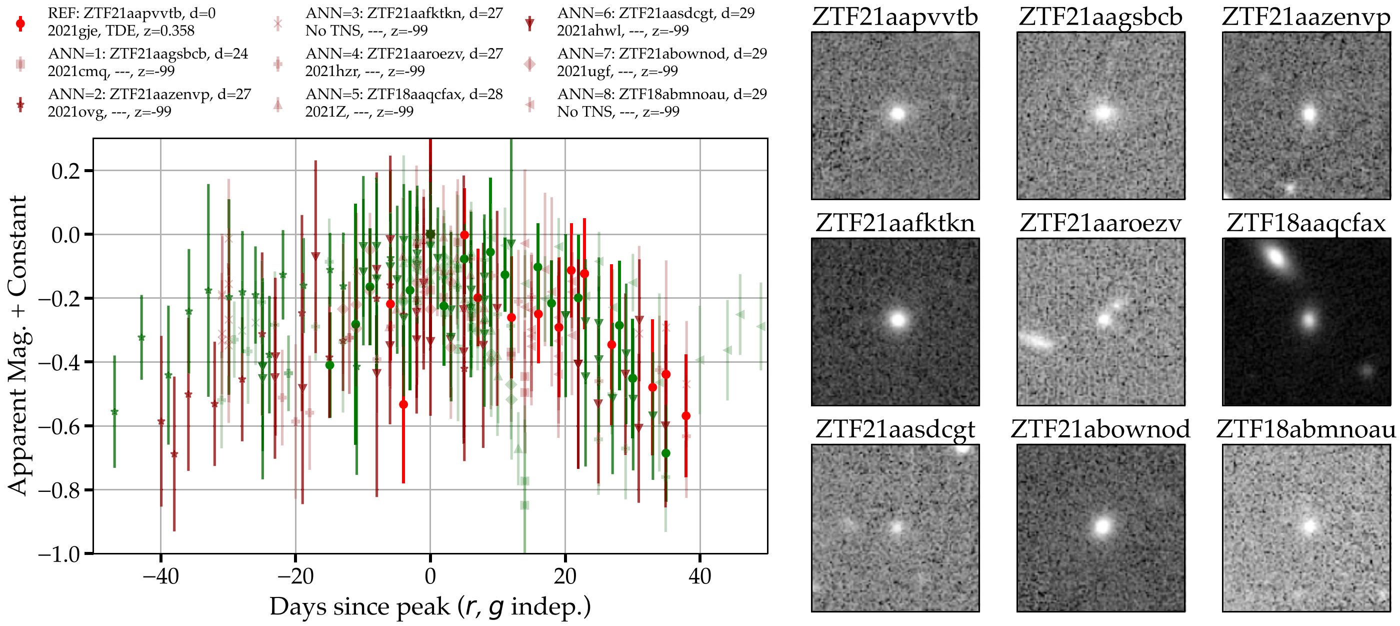

Figure 11. Same as Figures 9, 10 but for TDE 2021gje/ZTF21aapvvtb. In this case, ANN = 2 (AT 2021ovg/ZTF21aazenvp) was previously reported to TNS but never previously identified as a TDE candidate until this work (FLEET = 85% TDE), and ANN = 6 (AT 2021ahwl/ZTF21aasdcgt) is a missed TDE candidate (FLEET = 70% TDE; now reported to TNS from this work). We bold these new TDE candidates (dark green and dark red for ZTF-g and ZTF-r, respectively) for visual purposes. The similarity search in our database was completed in only ∼230 ms.

normal, an increase of 67%. This nearly doubles our total to 10 SN Ia-91bg-like objects.

5.3. Missed Opportunity SNe—Previously Reported to TNS

We can only reclassify SNe that have follow-up spectra. However, in principle, we can repeat the same ANN = 8 search of anomalous transients to look for potential anomalous candidates that were reported to TNS while active but have no classification spectrum. We call these "missed opportunity" transients. Although these events have faded, we identify unclassified ZTF transients with a high likelihood of being anomalous. While the entire search is ongoing and its entirety is beyond the scope of this work, we do identify a few candidates for which there is a strong possibility of being a TDE, using the FLEET-TDE classifier from Gomez et al. (2023d).

For a directed search, we examine the first eight ANNs of each of the 20 spectroscopic TDEs from our databank to look for missed opportunity TDEs that have yet to be reported in the literature. For instance, we use as reference 2021gje, a TDE at $z = 0.358$. We find the eight ANNs in ∼230 ms, as shown in Figure 11. Of those, two are AGNs, which have survived all our selection cuts (ANN = 3, ZTF21aafktkn; ANN = 8, ZTF18abmnoau), three are likely SN Ia (AT 2021cmq,[65] AT 2021Z, AT 2021ugf), one is an SLSN candidate[66] (AT 2021hzr), and two are likely TDEs or AGN flares with no previous reporting as such in the literature (AT 2021ovg, AT 2021ahwl). Moreover, AT 2021ahwl was never reported to TNS until this work (and thus was not considered in Gomez et al. 2023d).

AT 2021ovg has a $P(\mathrm{TDE}) = 85\%$ from FLEET from the late-time classifier, but there is no postpeak photometry, and thus, the usage of the late-time classifier is not fully appropriate here. From the FLEET-TDE classifier for prepeak, this object is assigned a $P(\mathrm{AGN}) = 60\%$ and $P(\mathrm{TDE}) = 32\%$. It is nuclear and blue, and there are no previous signs of variability according to ZTF DR data. The ALeRCE light-curve classifier (Sánchez-Sáez et al. 2021) assigns the highest probability as an SLSN at 34% (there is no TDE option), with the second highest probability being an AGN (16%). At an estimated redshift of $z = 0.241$ from SDSS (Csabai et al. 2003), we estimate an absolute magnitude of at least $M_{\mathrm{abs}} \approx -20.5$ mag, as there is no visible peak. AT 2021ovg was not reported in Gomez et al. (2023d).

AT 2021ahwl is a moderately strong missed TDE candidate at $P(\mathrm{TDE}) = 70\%$ from FLEET. It is nuclear, blue, and visible for ∼60 days, and there are no previous signs of strong variability according to ZTF DR data. At an estimated redshift of $z = 0.364$ from SDSS (Csabai et al. 2003), we estimate a peak absolute magnitude of $M_{\mathrm{abs}} \approx -21.8$ mag. Using the estimated redshift of $z = 0.372$ from PS1-STRM (Beck et al. 2021), we estimate a peak absolute magnitude of $M_{\mathrm{abs}} \approx -21.9$ mag. Both redshift estimates, if correct, are consistent with the reference z at 0.358.

We repeat a likewise examination for the 20 spectroscopic TDEs from our databank, and find one strong TDE candidate (AT 2021agpi with $P(\mathrm{TDE}) = 95\%$ from the FLEET late-time classifier), and three questionable ones (AT 2020hip, AT 2020yaf, AT 2021stx) that could be TDE or possibly AGN/AGN flares, or enhanced accretion-driven AGN flares (Trakhtenbrot et al. 2019). Spectroscopic observations are warranted for confirmation, although this opportunity has passed.

Similarly, because FLEET is also optimized for SLSN-I classification and discovery, we perform another directed search and examine the first eight ANNs of each of the 11 spectroscopic SLSN-I and 14 spectroscopic SLSN-II from our databank to look for missed opportunity SLSNe-I that have yet to be reported in the literature. We identify one moderately strong SLSN-I candidate (AT 2021fao with $P(\mathrm{SLSN-I}) = 64\%$ from the

---

[65] This object likely has a host misassociation, but is correct in the currently running `GHOST` version.

[66] AT 2021hzr has $P(\mathrm{SLSN-II}) = 68\%$ according to FLEET; though, FLEET is not optimized for SLSN-II classifications.

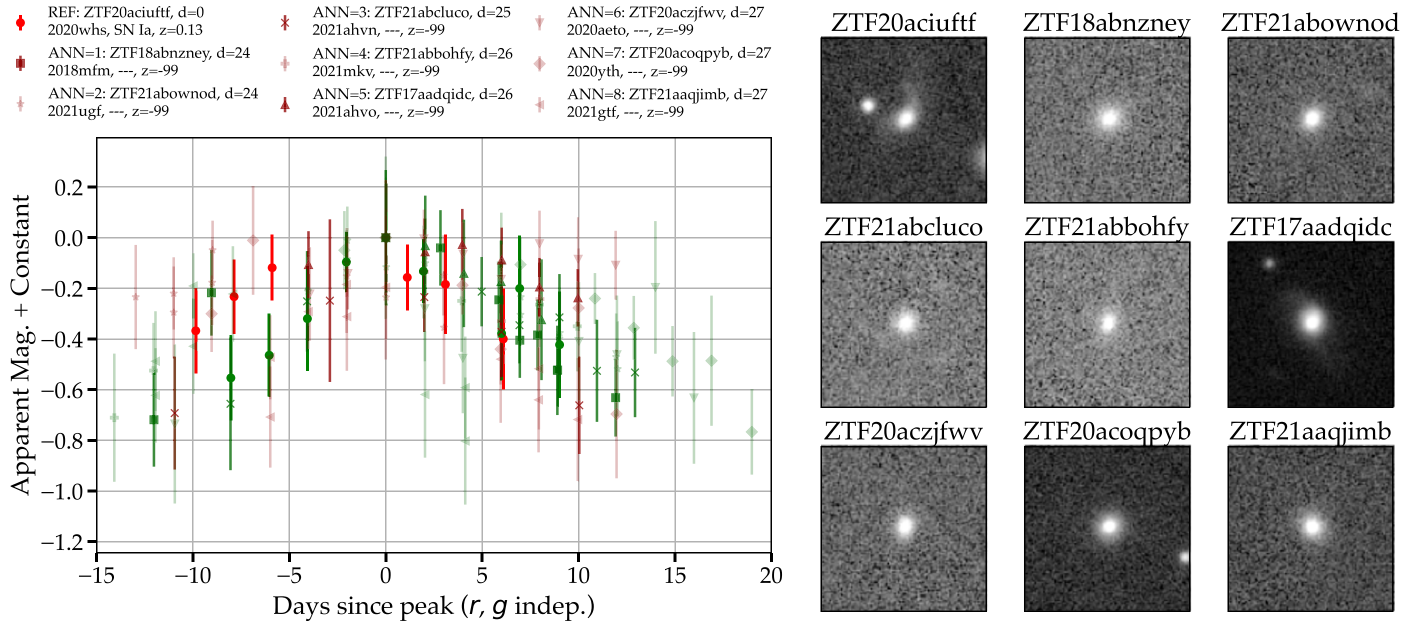

Figure 12. Same as Figures 9–11, but for Type Ia SN 2020whs/ZTF20aciuftf. In this case, ANN = 1 (AT 2018mfm/ZTF18abnzney), ANN = 3 (AT 2021ahvn/ZTF21abcluco), and ANN = 5 (AT 2021ahvo/ZTF17aadqidc) were missed SNe candidates (now reported to TNS from this work). We bold these SN discoveries (dark green and dark red for ZTF-g and ZTF-r, respectively) for visual purposes. The similarity search in our database was completed in only ∼140 ms.

FLEET late-time classifier) and one borderline candidate (AT 2021lnu with $P(\mathrm{SLSN-I}) = 49\%$ from the FLEET late-time classifier[67]).

5.4. Missed Opportunity SNe—Previously Unreported to TNS

Beyond identifying possible rare transients or anomaly candidates that were reported but have no spectroscopic follow-up, we can go the next step further and search for ANNs of known transients to find previously undiscovered or unreported transients (i.e., not reported to public catalogs like TNS). The idea here is that, if a known transient has a specific feature distribution, (approximate) nearest neighbors with similar distributions are likely to also be transients, and perhaps a fraction are not reported. An illustration of this directed search is found in Figure 12, whose search was completed in ∼140 ms.

The reference used is Type Ia SN 2020whs/ZTF20aciuftf at $z = 0.13$. From this ANN = 8 search, five were already reported to TNS, and thus, we manually investigated the remaining three with no submission. Upon expert analysis, we find that all 3 are missed SNe candidates: ANN = 1 (AT 2018mfm/ZTF18abnzney), ANN = 3 (AT 2021ahvn/ZTF21abcluco), and ANN = 5 (AT 2021ahvo/ZTF17aadqidc).

The first ANN, AT 2018mfm/ZTF18abnzney, is approximately 1.″38 from its host (WISEA J000044.86+152956.7). It peaks in ZTF-g at an apparent magnitude of 19.8 mag with a rise of 11 days, compared to the 19.6 mag peak apparent magnitude of reference SN 2020whs with an 8 day rise. At the time of the transient, AT 2018mfm has a peak brightness about 1 mag brighter than the baseline from ZTF DR5. Because of the short timescale and evolution, this is likely an SN Ia.

The third ANN, AT 2021ahvn/ZTF21abcluco, is remarkably similar in evolution to reference SN 2020whs. Both have a first epoch in ZTF-g at an apparent magnitude of 20.2 mag, rise to peak in 8 days to a peak apparent magnitude in ZTF-g of 19.6 mag, and fade in 12–13 days. This candidate has an estimated photo-z from from PS1-STRM (Beck et al. 2021) of $z = 0.168$, similar to the spectroscopic redshift of reference SN 2020whs ($z = 0.13$). This candidate, at approximately 0.″61 from the center of host WISEA J115609.31+210137.8, has an estimated peak absolute magnitude consistent with that of an SN Ia.

The fifth ANN, AT 2021ahvo/ZTF17aadqidc, is at the center (0.″17) of a galaxy (WISEA J100313.80+283944.8) with a spectroscopic redshift of 0.087, placing the peak $M_{\mathrm{abs}} \sim -18.6$ mag, well within the expected M_{abs} range of an SN I. Moreover, this object has seven ZTF-r and five ZTF-g alert stream observations and an apparent peak magnitude of $m \sim 19.4$ mag, which is about 1 mag brighter than the limiting magnitude of ZTF, making this a viable candidate for SN detection algorithms. A possible reason for why this was initially missed is due to its nuclear location, although this candidate's photometry is ∼0.3 mag above the baseline of from the ZTF DR and is red $g - r \approx 0.4$ mag (disfavoring an AGN flare). The ALeRCE light-curve classifier Sánchez-Sáez et al. (2021) predicts this candidate is an SN Ib/c.

We likewise perform a directed search by calculating the first eight ANNs of all spectroscopic transients in our databank, and investigate any nonreported ANNs for transient activity. In summary, we find that a fraction were never reported in public catalogs, totaling 84 transient candidate discoveries (not including likely AGN), all of which can be found in Table F1.

Finally, to be complete in recovering missed SNe from our databank, we perform an exhaustive search, visually inspecting each unique object not previously reported to TNS and not previously discovered through our ANN search. That search yielded 241 SNe candidates, as shown in Table F2.

In total, we report the discovery of 325 transients, all observed between 2018 and 2021 and absent from public catalogs (∼1% of all ZTF AT reports to TNS through 2021).

---

[67] Although not formally discussed in the literature, this object does have an Astronote identifying it as a ZTF SLSN candidate (Perley et al. 2021).

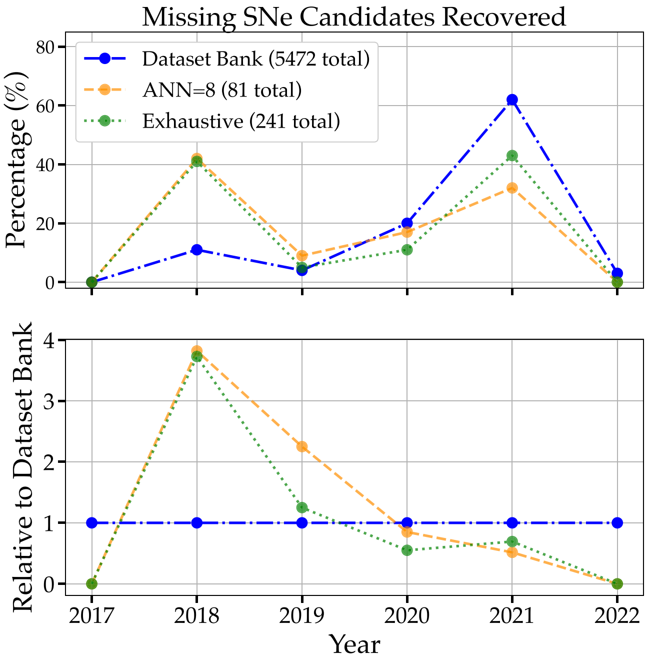

Figure 13. The unweighted percentage (top panel) and weighted (bottom panel) percentage relative to the fraction present in our original databank (blue) of the retroactive SNe candidate discoveries recovered by our ANN = 8 search method (gold) and exhaustive method (green). From both methods combined, we recover nearly equal numbers from 2018 and 2021 (134 out of 325 (41%) and 131 out of 325 (40%), respectively). However, relative to the yearly fractions that comprise our databank, we find that the fraction of our total 2018 discoveries relative to all discoveries from this work (41%) is nearly 3 times that of the fraction relative to that which comprise our databank (11%), whereas our total 2021 discoveries (40%) make up nearly two-thirds relative to our databank (62%). Overall, the weighted fraction of discoveries is greatest in 2018, the first full year of the ZTF survey, and decreasing every year through 2022, implying an improvement in detection and reporting methods via increasing numbers of active broker teams in the ZTF Alert Stream.

To better understand how our retroactive candidate SNe discoveries are distributed as a function of time, we plot in Figure 13 the unweighted percentage (top panel) and weighted (bottom panel) percentage relative to the fraction present in our original databank (blue) identified by our ANN = 8 search method (gold) and exhaustive method (green). The majority (134 out of 325, or 41%) of our candidate discoveries are from 2018 (with contributions of 35 out of 84 (42%) from our ANN = 8 search method and 99 out of 241 (41%) from our exhaustive search method). Contrary to our expectation, this is nearly equal to that of 2021, where we identify 131 out of 325 (also 40%) total missed candidates. However, our databank contains an uneven distribution of reported candidates per year. Grouping the IAU names, approximately 0% are from 2017, 11% are from 2018, 4% are from 2019, 20% are from 2020, 62% are from 2021, and 3% are from 2022. Thus, when we account for this unequal weighting, our result of 40% of our identified missed candidates from 2021 becomes unsurprising given that 62% of our databank is from 2021; it follows that there would be relatively more missed transients (numbers wise) for that year.

Note the weighted fraction of objects (bottom panel) missed in 2018 and 2019 is greater than that relative to 2020 and 2021. In fact, the weighted fraction is greatest in 2018, the first full year of the ZTF survey, and decreasing every year through 2022, implying an improvement in detection and reporting methods via increasing numbers of active broker teams in the ZTF Alert Stream.[68] For instance, the ALeRCE broker did not start real-time machine learning classification of the ZTF alert stream until early 2019 (Förster et al. 2021), and the ANTARES broker did not start reporting to TNS until 2021. Fink did not begin processing the ZTF public live-alert stream until 2019 November (Möller et al. 2021). Between additional teams reporting to TNS and the addition of object detection algorithms characterizing the ZTF alert stream in subsequent years (e.g., Muthukrishna et al. 2019; Andreoni et al. 2021; Carrasco-Davis et al. 2021; Coughlin et al. 2021; Duev & van der Walt 2021; Förster et al. 2021; Sánchez-Sáez et al. 2021; van Roestel et al. 2021; Aleo et al. 2022; Leoni et al. 2022; Reyes-Jainaga et al. 2023), it is sensible that we discover progressively fewer SNe candidates relative to those present in our databank. We encourage more retroactive study of the ZTF survey, particularly in 2018 and 2019 as these data are likely rife with undiscovered SNe and interesting candidates.

6. Discussion

6.1. Light-curve Extrapolation

Astronomical transients are the observational counterpart of terminal events, and a study of their photometric evolution can offer insights to underlying astrophysical properties. This can be performed on an individual basis (e.g., Gagliano et al. 2021), or across a population (e.g., Nyholm et al. 2020). Due to cadence, weather, and other factors intrinsic to astronomical observations, our photometry is irregularly sampled and noisy, and at times has large gaps spanning several epochs. Thus, obtaining homogeneous data suited for feature extraction and training machine learning and deep-learning models requires fast but accurate approximation. For such reasons, there are popular parametric light-curve models in the literature (Bazin et al. 2009; Villar et al. 2019; Russeil et al. 2024) and nonparametric models such as GP (Boone 2019; Demianenko et al. 2023) both of which predict the temporal flux evolution.

Note that although we make no such rigorous attempt at a formal light-curve fit model, we can approximate a possible light-curve evolution through considering the additional photometry and host galaxy environment from analogs to "reconstruct" a light curve. To our knowledge, this idea has not been demonstrated in the literature for light-curve fits, but similar ideas of applying "twin" SNe in regards to spectra for SN Ia standardization (Fakhouri et al. 2015; Boone et al. 2021) have shown promise.

Although we cannot recover photometry that was missed, we can place a soft constraint on the possible evolution by stacking light curves of ANNs, which broadly will be of a similar SN class, redshift, and host galaxy environment to the reference. Despite poor sampling of the light curve, or poor sampling at a critical phase, it is reasonable to investigate ANNs of these objects to gain insight of possible light-curve evolutionary paths and the associated variance. We show an example of reference SN Ia 2022cox/ZTF18aaiwewk in Figure 14, which is poorly sampled postpeak.

Despite having a large 30 day gap immediately after peak, we know from spectra SN 2022cox is an SN Ia at $z = 0.087$. Four of its ANN matches have spectroscopic classifications: three SN Ia with redshift ranging $z = 0.07$ to $z = 0.08$, and one poor SN II

---

[68] Only three objects in our databank have a ZTF objectID from 2017, so although it is included for completeness, it effectively bears no impact on this analysis.

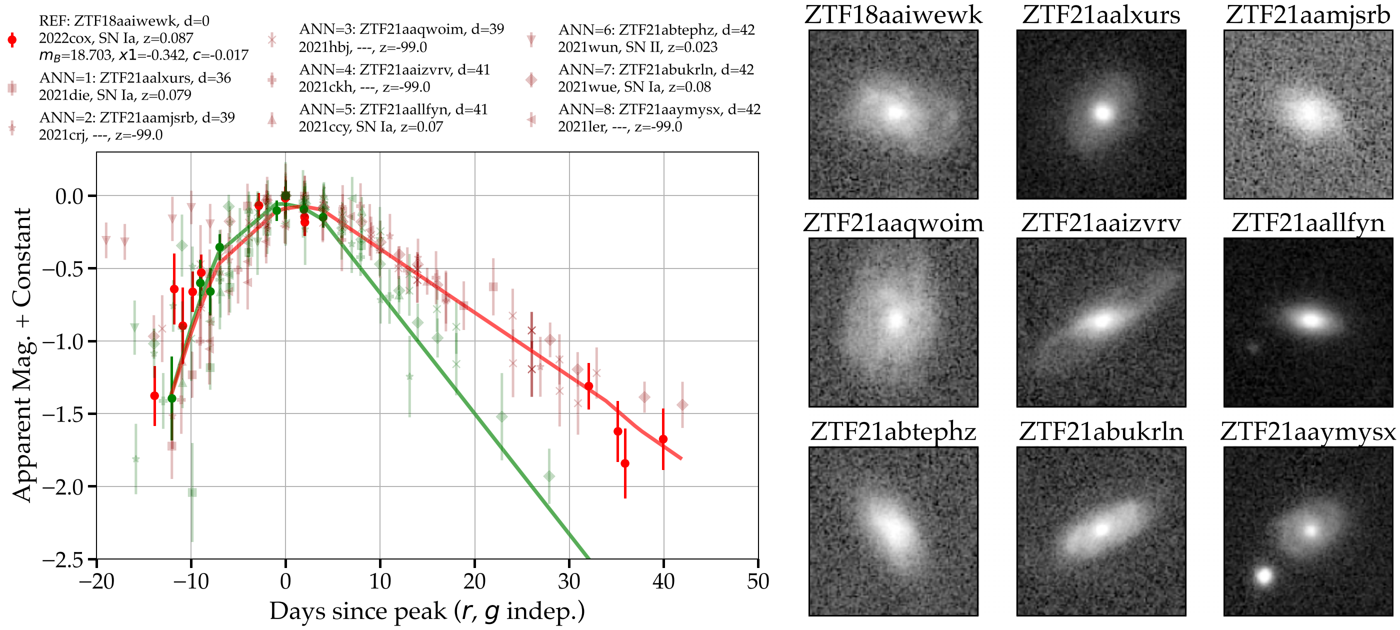

Figure 14. Same as Figures 9–12, but for Type Ia SN 2022cox/ZTF18aaiwewk, overplotted with the best-fit SALT3 model (Kenworthy et al. 2021). In this case, even though SN 2022cox is poorly sampled after peak with a 30 day gap, we can see fuller light-curve evolution and accompanying host galaxies of analogs through the similarity search, providing useful insights as to the possible classification and redshift range as well as the possible (unobserved) light-curve evolution of our reference. The SALT3 fit provides additional context, and we observe that analog light-curve epochs are positioned on or near this curve. The similarity search in our database was completed in only ∼2 ms.

match at $z = 0.023$. The photometric evolution of the other four ANNs appear to also be of an SN Ia nature. In total, only one appears to occur in an early-type galaxy. Stacking the light curves, the possible evolution postpeak (where the original gap is) is now well sampled. In fact, if we overplot the best-fit SALT3 model (Kenworthy et al. 2021), we find that the ZTF-r fit is well traced by the majority of the analogs' r-band photometry postpeak as well as the late-time photometry of our reference. Moreover, despite there being no ZTF-g photometry postpeak, the SALT3 fit is, too, well traced by the analogs' g-band photometry. This method may provide viable supporting evidence for light-curve constraints (or priors), particularly in cases where the SALT3 fit is poorly constrained.

The attraction of this approach is that this provides critical potential photometry where there is little to none, which contrasts nicely with a popular light-curve approximator: GP (Boone 2019; Demianenko et al. 2023; Gagliano et al. 2023). GP regression (Rasmussen & Williams 2005), while producing zero error at the learning points, incurs a large error in large gaps between subsequent epochs. A rigorous analytic comparison between the error of GP-fit light curves to those fit with analogs' photometry is left for future work.

Because our approach relies on the quality of the ANN search and consistency in light-curve profiles, it is necessary for there to be sufficient numbers of the SN type at varying redshifts and host galaxies to have requisite matches for tracing possible light-curve evolutions. Thus, for this work, it is unlikely this method will be effective for any SNe beyond normal SN Ia.

6.2. Common Failure Modes of Our Anomaly Detection Model

Despite the success of our anomaly detection model and our hyperparameter tuning to create a high fidelity sample, there are still cases where our model tags a nonanomaly as anomalous. Visual inspection of such objects from Tables 4–7, including their light-curve evolution and host galaxy association and environment, reveals some common patterns. The majority of misclassified anomalies tend to

1. have an incorrect host association (SN 2021aff, SN 2021uzt, SN 2021ojn, AT 2022ywi, SN 2023dgp, SN 2023eqx, AT 2023fli, AT 2023sws); this often occurs where the likely true host is faint and small (correlated with high-z); thus, normal light-curve behavior in one host galaxy environment could be exceedingly peculiar if wrongly associated to another vastly different host galaxy; however, we note that there are many instances where a spectroscopic or behavioral anomaly is correctly tagged anomalous despite a misassociation (SN 2018lnb, AT 2021dpa, AT 2023atr, SN 2021axu, SN 2023khp, SN 2023omf, AT 2023otw), potentially indicating that the object's light-curve behavior is peculiar enough alone to warrant being flagged anomalous;
2. be very blue (SALT3 $c \leqslant -0.2$) SN Ia or SN Ia candidates (AT 2020iga, AT 2020itp, AT 2020wwt, AT 2023tdy, AT 2023tim, AT 2023tjc, AT 2023tnr, AT 2023sws); this is likely a byproduct of `feature_linear_fit_slope_magn_g`) being our model's most discriminating feature for anomaly detection, where strong blue color is correlated with TDEs and some prepeak SNe IIn light curves; thus, it is unsurprising that SNe Ia significantly bluer than the SALT3 model are flagged anomalous;
3. be luminous (transient peak $M_{\rm abs} \leqslant -19.8$, using spectroscopic or photometric redshift estimates) Ia or Ia candidates whose host galaxy is faint (apparent brightness) and/or small (apparent size) and/or blue (likely spiral, star forming) according to the DESI Legacy Survey Imaging Surveys (AT 2020hjv, SN 2020qkx, SN 2020qql, AT 2020sgy, SN 2020wfg)[69];

---

[69] Note that SN Ia-CSM 2020kre is an example of successfully tagged event that falls in this category.

4. have no rise information (SN 2023cyx, SN 2023dgp, AT 2023fbj, AT 2023huz) or have no observations until the light curve is well in decline (SN 2020tnq, SN 2021vwx, SN 2023mcs); the latter possibly masquerades as a transient that is far fainter than is typical or is at a lower z than the host galaxy would imply, and would cause light-curve features that are OOD with respect to similar spectroscopic members due to the features' dependence on quantifying early (and intermediate) phases of light-curve behavior; in other words, two of the exact same SN Ia, where one is observed postpeak, and one that is observed for the entirety of the evolution will have some differing light-curve features despite being the same event;
5. be SN II with constant color evolution, thus have light curves that appear more SN Ib/c-like (SN 2019wmr, SN 2020acun); both examples were in YSE DR1 and misclassified by both ParSNIP and SuperRAENN as SN Ib/c, potentially indicating these events are in fact of the SNe IIb subtypes (the most common such case, as cited in Aleo et al. 2023);
6. have a significant (>10 day) gap in one passband during a period of intense color evolution of the light curve (SN 2023bfv, AT 2023glx); because ZTF-g and ZTF-r features are calculated independently, the passbands then no longer represent a similar phase in the evolution, but offset phases (where one lags behind), leading to behavior that is flagged as anomalous;
7. or have some combination of the above.

6.3. Host Galaxy SED Modeling with `PROSPECTOR`

As an additional analysis, we fit a spectral energy density model using stellar population synthesis to matched-host galaxies for each real-time anomalous transient candidate (i.e., tagged by `LAISS_RFC_AD_filter` *Filter*) to understand the connection between transients and their host galaxy environments. We perform Bayesian statistical inference to report the SFR (SFR [$M_\odot$ Gyr^{-1}]), the stellar mass of the galaxy (log(M)), the specific SFR (log(spSFR)), the age of the stellar population (age [Gyr]), and the V-band optical depth (A_V).

Our goal is to begin an investigation of the locations of tagged anomalies in host galaxy parameter feature-space (e.g., log(SFR) versus log(M)), relative to larger derived systematic samples across many normal and anomalous SNe classes from Schulze et al. (2021) and Sharma et al. (2023). We note that the sample from Schulze et al. (2021) contains the entirety of CC SNe from 2009 and 2017 by the Palomar Transient Factory—888 SNe of 12 distinct classes out to $z \approx 1$—including host galaxy photometric properties spanning the far-ultraviolet (FUV) to the mid-infrared (MIR) with modeled host galaxy SEDs to derive physical properties. Meanwhile, the sample from Sharma et al. (2023) added 12 Bright Transient Survey (BTS) SNe Ia-CSM hosts in comparison to BTS SNe Ia collected by Irani et al. (2022), with all such works overplotting hosts into SFR versus stellar mass feature-space. By comparing the locations of derived SFR and stellar mass for our tagged anomalies to larger systematic samples, we can incrementally add to the literature of derived host galaxy parameters for rare SNe, as well as investigate the nature of `LAISS`' tagged spectroscopically "Normal" SNe (e.g., SN Ia, SN II)—do they reside in sparse regions of host galaxy parameter space?

We analyze the set of matched-host galaxies from Table 5 using `PROSPECTOR`,[70] a framework integrating Bayesian statistical inference through nested (`DYNESTY`[71]) or distributed (`EMCEE`[72]) samplers with stellar population synthesis models (`FSPS`[73]; Foreman-Mackey et al. 2013; Leja et al. 2017; Johnson et al. 2021). We begin with the previously matched PanSTARRS host galaxies of anomaly candidate transients calculated via `GHOST` (Gagliano et al. 2021) within the `LAISS_RFC_AD_filter` *Filter*. Using the matched-host astrometry, we requery MAST[74] for PanSTARRS photometry (Chambers et al. 2016), choosing the Kron magnitude (Kron 1980) from the forced mean table in each band. The Kron magnitude utilizes an adaptive size aperture, the size of which we also query. Using the Kron aperture in the PanSTARRS r band, we query for the nearest fixed-size aperture and uncertainties in Galaxy Evolution Explorer (GALEX; Gezari et al. 2013), Two Micron All Sky Survey (2MASS; Skrutskie et al. 2006), AllWISE (Jarrett et al. 2011), and UKIRT Infrared Deep Sky Survey (UKIDSS; Lawrence et al. 2007) catalogs. Because the matched-host galaxies via `GHOST` are inherently from the PanSTARRS catalog, we are guaranteed to have PanSTARRS photometry.

To properly model the SED to extract physical and phenomenological parameters of the host galaxies, we require a vast range of host galaxy photometry from the FUV to the MIR. Thus, we only include tagged anomalies that have at least one observation in GALEX (filters FUV, NUV spanning 1542–2274 Å), PanSTARRS (filters g, r, i, z, y spanning 4776–9603 Å), either 2MASS (filters J, H, K_s spanning 12350–21590 Å) or UKIDSS (filters Z, Y, J, H, K spanning 8360–23800 Å), and AllWISE (filters W1, W2, W3, W4 spanning 33526–285500 Å). Moreover, we require a known redshift from spectroscopy, and the transient class (which is not an AGN) from at least one classification spectrum. This results in five matched hosts from the original sample of 45 tagged anomalies.

Our fitting routine follows that given in Leja et al. (2017); we first make a fit of our stellar population synthesis (SPS) model parameters using the Levenberg–Marquardt algorithm (Moré 1978) to initialize the Markov Chain Monte Carlo (MCMC) chains to a sensible parameterization. We then use the `Dynesty` sampler to perform Bayesian parameter estimation. Additional details on the SED fitting methods can be found in Appendix D.

Our reported galaxy properties are calculated from the parameter samples following Section 4.1 of Nugent et al. (2020), and we report the parameter estimates in Table 9. It is important to note that, while `PROSPECTOR` samples over the total solar masses formed, we estimate and report the stellar mass of the galaxy using the approximation given in Leja et al. (2013).

Of the five remaining objects after cuts, we only consider two bona fide anomalies, both of which are *spectroscopic* anomalies. However, we will find that all five objects lie in sparse regions of SFR versus stellar mass feature-space compared to Schulze et al. (2021) and Sharma et al. (2023), and could indicate that `LAISS`' tagged spectroscopically

---

[70] https://github.com/bd-j/prospector

[71] https://github.com/joshspeagle/dynesty

[72] https://github.com/dfm/emcee

[73] https://github.com/cconroy20/fsps

[74] https://mast.stsci.edu/portal/Mashup/Clients/Mast/Portal.html

Table 9
Host Galaxy Parameter Estimation with Prospector

ZTF ID	Spec. Class	Redshift	log(M)	Age (Gyr)	A_V	log(SFR)	log(spSFR)	χ^2/n.o.f.
ZTF23aatdcey	SN IIn	0.069	$10.128^{+0.100}_{-0.096}$	$4.406^{+1.952}_{-1.331}$	$1.636^{+0.267}_{-0.278}$	$-1.591^{+0.198}_{-0.238}$	$-11.719^{+0.241}_{-0.272}$	39.451/14
ZTF23aajkisd	SN IIn	0.029	$10.118^{+0.125}_{-0.110}$	$3.120^{+1.959}_{-1.031}$	$2.097^{+0.191}_{-0.184}$	$-1.469^{+0.167}_{-0.181}$	$-11.594^{+0.204}_{-0.209}$	23.446/14
ZTF23aapgswu	SN Ia	0.030	$8.367^{+0.077}_{-0.095}$	$10.401^{+2.368}_{-2.913}$	$0.117^{+0.189}_{-0.087}$	*$-4.389^{+0.570}_{-59.729}$*	*$-12.751^{+0.600}_{-59.850}$*	28.920/10
ZTF23aberpzw	SN II	0.024	$10.204^{+0.071}_{-0.083}$	$9.418^{+2.611}_{-2.387}$	$1.221^{+0.194}_{-0.188}$	$-1.949^{+0.175}_{-0.187}$	$-12.154^{+0.232}_{-0.216}$	19.036/14
ZTF23abayyjm	SN Ia	0.140	$10.815^{+0.110}_{-0.044}$	$1.471^{+1.006}_{-0.461}$	$3.275^{+0.184}_{-0.170}$	$0.147^{+0.263}_{-0.824}$	$-10.705^{+0.234}_{-0.761}$	28.451/11

Notes. We fit a spectral model using photometry from matched-host galaxies to each transient in Table 5 that pass strict selection cuts. We report the transient redshift, the log star formation rate (log(SFR) [$M_\odot$]), the total stellar mass of the galaxy (log(M) [$M_\odot$ yr^{-1}]), the log specific star formation rate (log(sSFR) [yr^{-1}]), the age of the stellar population (age (Gyr)) and the V-band optical depth (A_V). We also report the value of χ^2/n.o.f. (number of filters). For all distributions, we report the median values with upper and lower bounds equivalent to the 84th and 16th quantiles, respectively, unless otherwise noted (italicized text; see below). The italicized text indicates distributions that have unreliable median values (displayed as default in the table) due to their heavy negative skew. For these, the mode is better representative. For ZTF23aapgswu, the mode of log(SFR) is -3.998 $M_\odot$ yr^{-1} and the mode of log(spSFR) is -12.424 yr^{-1}. Bold text indicates an object we consider a likely anomaly.

"Normal" SNe have uncommon host galaxy parameters for their respective SNe classes despite small number statistics.

ZTF23aajkisd (SN IIn) at a log(M) of $10.118^{+0.1125}_{-0.110}$ $M_\odot$ and log(SFR) of $-1.469^{+0.167}_{-0.181}$ $M_\odot$ yr^{-1}, resulting in a log(spSFR) of $-11.594^{+0.204}_{-0.209}$ yr^{-1}, lies beyond the outer 90% contour of host galaxy populations in the mass–SFR plane as determined by a kernel density estimate of 111 SN IIn in the Schulze et al. (2021) sample (see their Figure 10). This indicates that SN 2023iex has an uncommonly low SFR for its relatively larger galaxy mass, resulting in a lower spSFR. However, in the mass–redshift plane (see their Figure 9), the host of SN 2023iex resides in the inner ~20% contour, indicating a common mass for its redshift (0.029).

ZTF23aatdcey (SN IIn) at a log(M) of $10.128^{+0.100}_{-0.096}$ $M_\odot$ and log(SFR) of $-1.591^{+0.198}_{-0.238}$ $M_\odot$ yr^{-1}, resulting in a log(spSFR) of $-11.719^{+0.241}_{-0.272}$ yr^{-1}, follows the same pattern as SN 2023iex (ZTF23aajkisd). The host of SN 2023nof (ZTF23aatdcey) also lies beyond the outer 90% contour of host galaxy populations in the mass–SFR plane from the Schulze et al. (2021) sample. Its host galaxy also resides in a sparsely populated, low region of spSFR, where the SFR is low considering the relatively larger mass of the galaxy. In the mass–redshift plane, it resides in the inner ~20% contour, indicating a common mass for its redshift (0.069).

ZTF23aberpzw (SN II), at a log(M) of $10.204^{+0.071}_{-0.083}$ $M_\odot$ and log(SFR) of $-1.949^{+0.175}_{-1.87}$ $M_\odot$ yr^{-1}, resulting in a log(spSFR) of $-12.154^{+0.232}_{-0.216}$ yr^{-1}, lies beyond the outer 90% contour of host galaxy populations in the mass–SFR plane as determined by a kernel density estimate of 498 SN II in the Schulze et al. (2021) sample (see their Figure 10). Moreover, it lies in what Schulze et al. (2021) found to be in the lower extreme of SFR—up to $\lesssim 2\%$ of regular CC SNe (SN Ibc, SN IIb, SN II, and SN IIn) whose hosts exhibit SFR $\lesssim 0.01$ $M_\odot$ yr^{-1} (a log(SFR) of -2 $M_\odot$ yr^{-1}) and spSFR between 10^{-14} yr^{-1} and 10^{-11} yr^{-1}, often with sizeable uncertainties. Thus, we consider this object to be a member of the lower extreme SFR hosts. In the mass–redshift plane, the host of SN 2023swf resides in the inner ~10% contour, indicating an extremely common mass at its redshift (0.024). Despite the light curve and spectra indicating a normal SN II, there is some evidence this object may be anomalous when considering the extremely low spSFR, and thus may be considered a *contextual* anomaly.

Overall, all CC SNe (ZTF23aajkisd, ZTF23aatdcey, ZTF23aberpzw) reside in a sparsely populated region of uncommonly low spSFR, where their galaxies have low SFR for their relatively larger size (all above the median and mode log(M) values as shown in Schulze et al. 2021, their Table 5).

ZTF23aapgswu (SN Ia), at a log(M) of $8.367^{+0.077}_{-0.095}$ $M_\odot$, is a peculiar case. Its host is a relatively compact, strongly blue, and likely high star-forming galaxy. One such consequence is a heavily negatively skewed SFR distribution, with one peak and long tail, resulting in a (more negative) median value of log(SFR) of $-4.389^{+0.570}_{-59.850}$ $M_\odot$ yr^{-1} and a log(spSFR) of $-12.751^{+0.600}_{-59.850}$ yr^{-1}. However, because of the heavy negative skew that biases the median, we also note a marginally more realistic (but still unphysical) mode value of log(SFR) of -3.998 $M_\odot$ yr^{-1} and a log(spSFR) of -12.424 yr^{-1}. Such an unphysically low SFR value can likely be traced to our choice of star formation history (SFH) function: it has been suggested that parameterized SFH are not flexible enough to deal with the actual complexities inherent in galaxy systems (Leja et al. 2019a). As a further point of caution, no SN Ia from Irani et al. (2022) of BTS SNe Ia have such a low log(SFR), and thus, we make no strong claims as to the validity of the fit. More investigation of the host galaxy and derived host parameters of SN 2023mcs is encouraged, especially when considering it is host to an SN Ia.

ZTF23abayyjm (SN Ia) at a log(M) of $10.815^{+0.110}_{-0.044}$ $M_\odot$ (the largest in our sample of five objects) and log(SFR) of $0.147^{+0.263}_{-0.842}$ $M_\odot$ yr^{-1}, resulting in a log(spSFR) of $-10.705^{+0.234}_{-0.761}$ yr^{-1}, is the only event that resides in a host galaxy with above Milky Way metallicity ([Z/H] = 0.0; see Choi et al. 2016). When compared to BTS SNe Ia from Irani et al. (2022) and Sharma et al. (2023), it is at the upper regime of massive, highly star-forming, and high metallicity host galaxies beyond the red and blue contours corresponding to Galaxy Zoo ellipticals and spirals (see Figure 14 of Irani et al. 2022). However, the spSFR is average after normalizing SFR by the very large galaxy mass, and is among the highest mass hosts compared to the Irani et al. (2022) and Sharma et al. (2023) samples.

The goal of this exercise was to lay the groundwork of building a downstream *Filter* to fit SPS host galaxy models. In this pursuit, we have discerned valuable insights. Moreover, we recognize the need for additional technologies and software infrastructure that needs to be developed before this process can be scaled to *Rubin*-era data streams.

First, it was critical to use matched-aperture size photometry to get sensible results. Even so, we chose to only use precalculated magnitudes from reported fixed-size apertures, when in principle the apertures should be matched exactly across each band. Additional work would be necessary to create a robust pipeline that works on galaxies at various distances. Our fits could likely be improved by utilizing nonparameteric SFH functions. Additionally, MCMC sampling is computationally expensive. Each object's sampling takes approximately 2 hr on a shared compute node equipped with an AMD EPYC 7502 CPU. To rectify this, other works have suggested speeding up SPS models with neural network emulators (Kwon et al. 2023), which are reported to achieve a 100× speed-up; or with approximate posterior distributions of galaxy parameters efficiently via simulation-based inference (Cranmer et al. 2020) and amortized neural posterior estimation (Papamakarios & Murray 2016) techniques, which can achieve an amortized calculation of galaxy parameter posteriors in less than 1 s (see Hahn & Melchior 2022; Khullar et al. 2022; Wang et al. 2023, and references therein); or with a physics-informed VAE (Gagliano & Villar 2023). These efforts are promising and ongoing, but beyond the scope of this work.

7. Conclusion

In this work, we present `LAISS`, a pipeline for real-time anomaly detection and approximate similarity searches of astronomical transients within large volumetric data streams. We debut our anomaly detection model as a *Filter* on the ANTARES broker to process the nightly ZTF Alert Stream. Our model is based on an RFC architecture using extracted light-curve and contextual host galaxy features without the need for redshift information, designed to classify several types of transients as anomalies ("Anomaly") with high purity:

1. spectroscopic anomalies (i.e., designate all SNe other than normal Type Ia, Type Ia-91T-like, and normal Type II, IIP as anomalous);
2. contextual anomalies (e.g., an SN in an atypical galactic environment);
3. behavioral anomalies (e.g., an SN rebrightening due to CSM-interaction).

We run our anomaly detection model on the active ZTF alert stream for real-time discovery, as well as legacy subsets of the ZTF alert stream and the YSE DR1 for retroactive discovery. Moreover, we construct a low-latency approximate similarity search model within our derived light-curve and host galaxy feature-space to find transient analogs of similar light-curve evolution and host galaxy environments.

Our conclusions and key takeaways for the `LAISS` anomaly detection and approximate similarity search pipeline are as follows:

1. The most important *light-curve features* for distinguishing anomalies from other transients according to our RFC model using both impurity and permutation importances are the light-curve slope in a least squares fit of the linear stochastic model with Gaussian noise described by observation errors $\{\delta_i\}$ (`feature_linear_fit_slope_magn_g`) and its error in r band (`feature_linear_fit_slope_sigma_magn_r`), and the unbiased Anderson–Darling normality test statistic for g-band flux (`feature_anderson_darling_normal_flux_g`).
2. The most important *host galaxy features* for distinguishing anomalies from other transients according to our RFC model using both impurity and permutation importances are radial offset and color-derived features, namely, the aperture magnitude $i-z$ color difference (`i-z`) and the normalized host separation `dist/DLR`. Notably, the high importance of these contextual host galaxy features is in agreement with that of FLEET (Gomez et al. 2020, 2023a, 2023d) and `GHOST` (Gagliano et al. 2021), despite the nature of their different anomaly detection and SNe classification tasks.
3. Using the spectroscopic label alone from TNS (before vetting) as an indicator of an anomalous transient, we can recover anomalies with a purity of ≈50% using the full light-curve classifier for events with $P(\text{anom}) > 0.5$. If we assume anomalies (defined in this work as everything not Type Ia-normal, 91T-like, II-normal, and IIP) comprise about 10% of the SNe population in a magnitude-limited survey like ZTF, this is a factor of ~5 improvement compared to random selection. The corresponding completeness for this threshold is 29%. However, if we additionally consider *updated spectroscopic classifications*, *contextual anomalies*, and *behavioral anomalies* (such as SNe in an atypical environment or peculiar light-curve or spectral behavior) as a result of expert vetting, we can recover anomalies with an *observed purity* of ≈90% using the full light-curve classifier for events with $P(\text{anom}) > 0.5$ at a completeness threshold of 38%. This is a factor of ~8 improvement compared to random selection.
4. When using any SNe as input into our approximate similarity search, we can retrieve eight ANNs in our 5472 object, 60-dimensional principal component-space in $\mathcal{O}(100)$ ms. If needed, we can reextract light-curve features in ~1 s and host galaxy association and features in ~1 minute. We show that an ANN solution to finding SNe analogs with precomputed features can scale to *Rubin* data streams.
5. We demonstrate an approximate similarity search for finding SNe analogs is useful in many applications, such as but not limited to the following: finding similar SNe evolution and host galaxy environment, finding missed SNe in legacy data sets, providing possible light-curve evolution of poorly sampled SNe, and prompting some reclassification of SNe (often to a rarer subclass of the SNe type).
6. From our approximate similarity search, we report 17 unique reclassified SNe and 84 previously undiscovered SNe candidates found by an ANN = 8 search of our databank spectroscopic sample. From a final exhaustive search of our databank, we report the remaining 241 SNe candidates to TNS. In total, we report 325 discoveries, all from 2018 to 2021 and absent from public catalogs (~1% of all ZTF AT reports to TNS through 2021[75]).
7. Despite small number statistics, there is some evidence to suggest that objects tagged anomalous by `LAISS` have host galaxy properties, such as mass, SFR, and specific

---

[75] https://tinyurl.com/ZTF-AT-TNS-reports-thru2021

SFR, that reside in sparse regions of latent space of equivalent SN type when compared to larger samples in the literature.

7.1. List of Anomalies

We report the following new *spectroscopic* and *behavioral* anomalies, either flagged in real-time or retrospectively by our anomaly detection model. They are confirmed via follow-up observations or from a retroactive reclassification of existing spectra prompted by approximate similarity search. Additionally, we report likely anomalous transient candidates based on their light-curve evolution, host galaxy environment, and results from photometric classifiers such as FLEET (Gomez et al. 2020). We newly identify a myriad of peculiar and rare transients, including the following:

1. *SLSN (11 total; 2 spectroscopic; 9 photometric)*. We identify 11 new candidate members, 2 spectroscopic (ZTF21aaizyqc/SN 2021ckb, ZTF22abfdzrv/SN 2022vmg); 9 photometric (ZTF20acbiwfi/AT 2020afew, ZTF20aczhbki/AT 2020afex, ZTF21aajrgkw/AT2021dpa, ZTF21aaualdi/AT 2021ahwa, ZTF21abasbpe/AT 2021ahwh, ZTF20acyroio/AT 2021ahwp, ZTF23aahjdxa/AT 2023gpp, ZTF23aaqqeek/AT 2023mne, ZTF23aawblmi/AT 2023otw).
2. *TDE (10 total; 1 spectroscopic, 9 photometric)*. We identify nine new candidate members, one spectroscopic (ZTF23aamsetv/SN 2023kvy); eight photometric (ZTF18abtjrbt/AT 2018mfz, ZTF18acvwkvc/AT 2018mkd, PS19gzf/AT 2019vuz, ZTF20acpgnmw/AT 2020afev, ZTF20acpzjsk/AT 2020affx, ZTF21aazenvp/AT 2021ovg, ZTF21aasdcgt/AT 2021ahwl, ZTF22absuavp/AT 2022zyh, ZTF22abzajwl/AT 2023adr).
3. *Long-rising (>40 day) SN II (six total; two spectroscopic, four photometric)*. We identify six new candidate members, two spectroscopic (ZTF20abqlmwn/SN 2020rmk, ZTF21aaqyifh/SN 2021hqe); four photometric (ZTF18abiitmq/AT 2018mhh, ZTF18abkmfaj/AT 2018mhj, ZTF21abasjcd/AT 2021rmq, ZTF23aajestr/AT 2023inr).
4. *SN Ic-BL (one total; one spectroscopic, zero photometric)*. We identify one new candidate member; one spectroscopic (ZTF21aacufip/SN 2021vz); zero photometric.
5. *SN Ib (two total; two spectroscopic, zero photometric)*. We identify two new candidate members, two spectroscopic (ZTF21aaqwfqe/SN 2021hen, ZTF21aabyifm/SN 2021qv); zero photometric.
6. *SN Ib/c (six total; zero spectroscopic, six photometric)*. We identify six new candidate members, zero spectroscopic; six photometric (ZTF18abwhsnx/AT 2018mgd, ZTF18aajgowk/AT 2018mgw, ZTF20acuyxki/AT 2020afes, ZTF21aaiapis/AT 2021ahwg, ZTF23aafgmaz/AT 2023frg, ZTF23abedgfr/AT 2023syt).
7. *SN IIb (four total; three spectroscopic, one photometric)*. We identify four new candidate members, three spectroscopic (ZTF21aajvukh/SN 2021cpi, ZTF21abnvlnj/SN 2021tyf, ZTF23aarktow/SN 2023myo); one photometric (ZTF18acvgmpx/AT 2018mkc).
8. *SN IIn (19 total; 4 spectroscopic, 15 photometric)*. We identify 19 new candidate members, 4 spectroscopic (ZTF21abcjpnm/SN 2021njo, ZTF23aatdcey/SN 2023nof, ZTF23aatcsou/SN 2023nwe, ZTF23aavtugd/SN 2023omf); 15 photometric (ZTF18ablqjws/AT 2018mga, PS20czv/AT 2020jvi, PS20mgo/AT 2020acyu, ZTF20aclghmy/AT 2020affa, ZTF20acbptqx/2020affg, ZTF20acngnvf/AT 2020affn, ZTF20acrssoi/AT 2020affz, ZTF20actkqax/AT 2020afgb, ZTF20acxyrkk/AT 2020afge, ZTF21abiggqx/AT 2021rjf, ZTF22aatwxrl/AT 2022oym, ZTF23aaahnss/AT 2023atr, ZTF23aaveoxd/AT 2023ofr, ZTF23aaufkak/AT 2023nwk, ZTF23aaewyhm/AT 2023gzn).
9. *SN Ia-CSM (three total; two spectroscopic, one photometric)*. We identify three new candidate members, two spectroscopic (ZTF20acpbboa/SN 2020ywo, ZTF23aatabje/SN 2023ocx); one photometric (ZTF23aaqbyzr/SN 2023mic).
10. *SN Ia-91bg-like (10 total; 7 spectroscopic, 3 photometric)*. We identify 10 new candidate members, 7 spectroscopic (ZTF20acnznol/SN 2020yje, ZTF21abmwgow/SN 2021ttg, ZTF21acfigoo/SN 2021aazj, ZTF21achjwus/SN 2021abpz, ZTF21acjgafq/SN 2021ackd, ZTF21acmnpqa/SN 2021adpx, ZTF23abhafym/SN 2023tsw); 3 photometric (ZTF19aavoqbe/AT 2019aatq, ZTF21aaiahsu/AT 2021ahxg, ZTF23aaflptz/AT 2023gbk).
11. *Unknown blue transients (two total; two photometric)*. We identify two blue transients of unknown type, zero spectroscopic; two photometric (ZTF21aabyoxk/AT 2021ahwu, ZTF21aawzmne/AT 2021ahyj).

Lastly, we aggregate the *contextual* anomalies:

1. *SN candidates in rare ring host galaxy (2)*. This includes ZTF18acvgmpx/AT 2018mkc and ZTF21aakjxhg/AT 2021ahxn.
2. *SN Ia in flocculent spiral/merger (1)*. This includes ZTF20aazpwen/SN 2020kbl.
3. *SN II or SN II candidate in possible elliptical host galaxy (3)*. This includes ZTF19abljudj/AT 2019aats, ZTF20aawlmfu/AT 2020kmj, and ZTF23abcinum/SN 2023sds.

Acknowledgments

We thank Anya Nugent, Nell Byler, Maggie Verrico, Nicholas Earl, Yuxin Dong, Charlie Conroy, and Ben Johnson for helpful correspondence regarding host galaxy SED modeling.

Author contributions are listed below.

P.D.A. as project lead and YSE collaboration meeting co-lead contributed to ideation, development, and deployment of `LAISS`; statistical and data analysis; analysis of classified objects; lead writing and lead editing; and figures.

A.E. contributed to writing text, making figures, and performing the analysis using `Prospector`.

G.N. as YSE Executive Committee member contributed to oversight; editing; and helpful discussions.

C.R.A. contributed to observing and reducing Nordic Optical Telescope (NOT) data used for classification; analysis of classified objects; and sniffing and/or sorting and/or flagging transients.

K.M. contributed to ANTARES operations (software); *Filter* management; and draft review.

K.A. contributed to WiFeS observations to confirm candidates; and to draft review.

V.F.B. as YSE Executive Committee member contributed to helpful discussions.

A.B. contributed to the creation of the Slack-bot.

T.J.L.d.B. contributed to PS1 operations.

B.M.B. contributed as a YSE junior review panelist.

K.C.C. contributed to PS1 operations.

K.W.D. contributed to observing and reducing Lick 3 m data used for classification.

N.E. contributed to ANTARES operations (software); and ANTARES operations (hardware).

D.F. contributed as a YSE junior review panelist.

R.J.F. as YSE Executive Committee member contributed to helpful discussions, managed the YSE survey, and was PI of spectroscopic observing programs.

A.T.G. contributed to draft review; helpful discussions; and sniffing and/or sorting and/or flagging transients.

C.G. contributed to sniffing and/or sorting and/or flagging transients. H. Gao contributed to PS1 operations.

S.G. contributed to observing and reducing SOAR data used for classification in coordination with the FLEET program; and to helpful discussions.

M.G. contributed as lead YSE junior review panelist.

D.O.J. contributed to helpful discussions and to PS1 operations.

C.-C.L. contributed to PS1 operations.

E.A.M. contributed to PS1 operations.

K.S.M. as YSE Executive Committee contact contributed to oversight and helpful discussions.

T.M. contributed to ANTARES operations (software); discussions; and draft review.

S.I.R. contributed to sniffing and/or sorting and/or flagging transients.

V.G.S. contributed as a YSE junior review panelist.

M.D.S. contributed to ANTARES operations (software); discussions; and draft review.

K.M.d.S. contributed as a YSE junior review panelist.

S.V. contributed to ANTARES operations (software); and *Filter* management.

V.A.V. as YSE Executive Committee member contributed to helpful discussions.

R.J.W. contributed to PS1 operations.

P.D.A.'s and G.N.'s contributions to this work were directly supported by NSF AST–2206195, and we gratefully acknowledge this funding. P.D.A. has also been supported by the Illinois Survey Science Graduate Fellowship from the Center for AstroPhysical Surveys (CAPS; https://caps.ncsa.illinois.edu/) at the National Center for Supercomputing Applications (NCSA). A.W.E. was partially supported by the Open Call Initiative, under the Laboratory Directed Research and Development (LDRD) Program at Pacific Northwest National Laboratory (PNNL). PNNL is a multiprogram national laboratory operated for the U.S. Department of Energy (DOE) by Battelle Memorial Institute under contract No. DE-AC05-76RLO 1830. G.N. is also supported by NSF CAREER grant AST–2239364, funded in-part by a grant by Charles Simonyi, and NSF OAC–2311355, DOE support through the Department of Physics at the University of Illinois, Urbana-Champaign (No. 13771275), and support from the HST Guest Observer Program through HST-GO-16764 and HST-GO-17128 (PI: R. Foley). Support was provided by Schmidt Sciences, LLC. for K.M. V.F.B.'s contributions to this work were supported by NSF AST-2206165. B.M.B. is supported by the Cambridge Centre for Doctoral Training in Data-Intensive Science funded by the UK Science and Technology Facilities Council (STFC). The UCSC team is supported in part by NASA grant NNG17PX03C, NSF grants AST–1815935 and AST–2307710, the Gordon & Betty Moore Foundation, the Heising-Simons Foundation, and a fellowship from the David and Lucile Packard Foundation to R.J.F. A.T.G. is supported by the National Science Foundation under cooperative agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, http://iaifi.org/). C.G. is supported by a VILLUM FONDEN Young Investigator Grant (project No. 25501). M.G. and K.S.M. are supported by the European Union's Horizon 2020 research and innovation program under ERC grant agreement No. 101002652 and Marie Skłodowska-Curie grant agreement No. 873089. K.M.d.S. acknowledges support by the NSF through grant AST–2108676. K.d.S. thanks the LSST-DA Data Science Fellowship Program, which is funded by LSST-DA, the Brinson Foundation, and the Moore Foundation; her participation in the program has benefited this work. V.S. acknowledges the support of the LSST Corporation's 2021 Enabling Science award for undergraduates. V.A.V. acknowledges support by the NSF through grant AST–2108676.

The Young Supernova Experiment (YSE) and its research infrastructure is supported by the European Research Council under the European Union's Horizon 2020 research and innovation program (ERC grant agreement 101002652, PI K. Mandel), the Heising-Simons Foundation (2018-0913, PI R. Foley; 2018-0911, PI R. Margutti), NASA (NNG17PX03C, PI R. Foley), NSF (AST–1720756, AST–1815935, AST–2307710, PI R. Foley; AST-1909796, AST-1944985, PI R. Margutti), the David & Lucille Packard Foundation (PI R. Foley), VILLUM FONDEN (project 16599, PI J. Hjorth), and the Center for AstroPhysical Surveys (CAPS) at the National Center for Supercomputing Applications (NCSA) and the University of Illinois Urbana-Champaign.

The ANTARES project has been supported by the National Science Foundation through a cooperative agreement with the Association of Universities for Research in Astronomy (AURA) for the operation of NOIRLab, through an NSF INSPIRE grant to the University of Arizona (CISE AST-1344024, PI: R. Snodgrass), and through a grant from the Heising-Simons Foundation.

YSE-PZ was developed by the UC Santa Cruz Transients Team. The UCSC team is supported in part by NASA grants NNG17PX03C, 80NSSC18K0303, 80NSSC19K0113, 80NSSC19K1386, 80NSSC20K0953, 80NSSC21K2076, 80NSSC22K1513, 80NSSC22K1518, and 80NSSC23K0301; NSF grants AST–1720756, AST–1815935, and AST–1911206; grants associated with Hubble Space Telescope programs DD–14925, DD–15600, GO–15876, GO–16238, SNAP–16239, GO–16690, SNAP–16691, and GO–17128; the Gordon & Betty Moore Foundation; the Heising-Simons Foundation; fellowships from the Alfred P. Sloan Foundation and the David and Lucile Packard Foundation to R.J.F.; Gordon and Betty Moore Foundation postdoctoral fellowships and a NASA Einstein fellowship, as administered through the NASA Hubble

Fellowship program and grant HST-HF2-51462.001, to D.O.J.; and a National Science Foundation Graduate Research Fellowship, administered through grant No. DGE-1339067, to D.A.C.

Pan-STARRS is a project of the Institute for Astronomy of the University of Hawaii, and is supported by the NASA SSO Near Earth Observation Program under grants 80NSSC18K0971, NNX14AM74G, NNX12AR65G, NNX13AQ47G, NNX08AR22G, 80NSSC21K1572 and by the State of Hawaii. The PS1 Surveys and the PS1 public science archive have been made possible through contributions by the Institute for Astronomy, the University of Hawaii, the Pan-STARRS Project Office, the Max-Planck Society and its participating institutes, the Max Planck Institute for Astronomy, Heidelberg and the Max Planck Institute for Extraterrestrial Physics, Garching, The Johns Hopkins University, Durham University, the University of Edinburgh, the Queen's University Belfast, the Harvard-Smithsonian Center for Astrophysics, the Las Cumbres Observatory Global Telescope Network Incorporated, the National Central University of Taiwan, STScI, NASA under grant NNX08AR22G issued through the Planetary Science Division of the NASA Science Mission Directorate, NSF grant AST-1238877, the University of Maryland, Eotvos Lorand University (ELTE), the Los Alamos National Laboratory, and the Gordon and Betty Moore Foundation.

Parts of this work are based on observations obtained with the Samuel Oschin Telescope 48 inch and the 60 inch Telescope at the Palomar Observatory as part of the Zwicky Transient Facility project. ZTF is supported by the National Science Foundation under grant Nos. AST–1440341 and AST–2034437 and a collaboration including current partners Caltech, IPAC, the Weizmann Institute of Science, the Oskar Klein Center at Stockholm University, the University of Maryland, Deutsches Elektronen-Synchrotron and Humboldt University, the TANGO Consortium of Taiwan, the University of Wisconsin at Milwaukee, Trinity College Dublin, Lawrence Livermore National Laboratories, IN2P3, University of Warwick, Ruhr University Bochum, Northwestern University and former partners the University of Washington, Los Alamos National Laboratories, and Lawrence Berkeley National Laboratories. Operations are conducted by COO, IPAC, and UW. The ZTF forced-photometry service was funded under the Heising-Simons Foundation grant No. 12540303 (PI: Graham).

Parts of this research were supported by the Australian Research Council Discovery Early Career Researcher Award (DECRA) through project No. DE230101069.

A major upgrade of the Kast spectrograph on the Shane 3 m telescope at Lick Observatory was made possible through generous gifts from the Heising-Simons Foundation as well as William and Marina Kast. Research at Lick Observatory is partially supported by a generous gift from Google.

This research has made use of the VizieR catalog access tool, CDS, Strasbourg, France (10.26093/cds/vizier). The original description of the VizieR service was published in Ochsenbein et al. (2000).

Facilities: ADS, TNS, NED, ANTARES.

Software: `ANNOY` (Bernhardsson 2018), `antares-client` (NOIRLab CSDC Team 2023), `Astropy` (Astropy Collaboration et al. 2013; Price-Whelan et al. 2018), `FLEET` (Gomez et al. 2020), `GHOST` (Gagliano et al. 2021), `Matplotlib` (Hunter 2007), `numpy` (van der Walt et al. 2011), `Pandas` (The pandas development team 2020), `Scikit-Learn` (Pedregosa et al. 2011), `YSE-PZ` (Coulter et al. 2022, 2023).

Appendix A
Light-curve Features

Our light-curve features are extracted with the `lc_feature_extractor`filter in ANTARES using the `light-curve` package. We extract the same features for r and g band, comprising 62 total features (31 r band, 31 g band). A brief description of each feature is as follows:

1. `feature_amplitude_magn`, half amplitude of magnitude (Malanchev et al. 2021);
2. `feature_anderson_darling_normal_magn`, unbiased Anderson–Darling normality test statistic for magnitude (Malanchev et al. 2021);
3. `feature_beyond_1_std_magn`, fraction of observations beyond $n = 1\sigma_m$ from the mean magnitude $\langle m \rangle$ (D'Isanto et al. 2016);
4. `feature_beyond_2_std_magn`, fraction of observations beyond $n = 2\sigma_m$ from the mean magnitude $\langle m \rangle$ (D'Isanto et al. 2016);
5. `feature_cusum_magn`, a range of cumulative sums dependent on the number of observations, mean magnitude, and magnitude standard deviation (Kim et al. 2014);
6. `feature_inter_percentile_range_2_magn`, interpercentile range for $p = 0.02$, where p is the pth quantile of the magnitude distribution (Malanchev et al. 2021);
7. `feature_inter_percentile_range_10_magn`, interpercentile range for $p = 0.10$, where p is the pth quantile of the magnitude distribution; a special case of the interpercentile range known as the interdecile range (Malanchev et al. 2021);
8. `feature_inter_percentile_range_25_magn`, interpercentile range for $p = 0.25$, where p is the pth quantile of the magnitude distribution; a special case of the interpercetile range known as the interquartile range (Malanchev et al. 2021);
9. `feature_kurtosis_magn`, excess kurtosis of magnitude (Malanchev et al. 2021);
10. `feature_linear_fit_slope_magn`, the slope of the light curve in the least squares fit of the linear stochastic model with Gaussian noise described by observation errors $\{\delta_i\}$ (Malanchev et al. 2021);
11. `feature_linear_fit_slope_sigma_magn`, the error of the slope of the light curve in the least squares fit of the linear stochastic model with Gaussian noise described by observation errors $\{\delta_i\}$ (Malanchev et al. 2021);
12. `feature_magnitude_percentage_ratio_40_5_magn`, the magnitude 40 to 5 ratio, written in terms of the magnitude distribution quantile function Q. (D'Isanto et al. 2016);
13. `feature_magnitude_percentage_ratio_20_5_magn`, the magnitude 20 to 5 ratio, written in terms of the magnitude distribution quantile function Q. (D'Isanto et al. 2016);
14. `feature_mean_magn`, the nonweighted mean magnitude;

15. `feature_median_absolute_deviation_magn`, the median of the absolute value of the difference between magnitude and its median (D'Isanto et al. 2016);
16. `feature_percent_amplitude_magn`, the maximum deviation of magnitude from its median (D'Isanto et al. 2016);
17. `feature_median_buffer_range_percentage_10_magn`, the fraction of observations inside Median $(m) \pm 10 \times (\max(m) - \min(m))/2$ interval (D'Isanto et al. 2016);
18. `feature_median_buffer_range_percentage_20_magn`, the fraction of observations inside Median $(m) \pm 20 \times (\max(m) - \min(m))/2$ interval (D'Isanto et al. 2016);
19. `feature_percent_difference_magnitude_percentile_5_magn`, ratio of $p = $ fifth interpercentile range to the median (Malanchev et al. 2021);
20. `feature_percent_difference_magnitude_percentile_10_magn`, ratio of $p = $ 10th interpercentile range to the median (Malanchev et al. 2021);
21. `feature_skew_magn`, skewness of magnitude, G_1 (Malanchev et al. 2021);
22. `feature_standard_deviation_magn`, standard deviation of magnitude, σ_m (Malanchev et al. 2021);
23. `feature_stetson_k_magn`, Stetson K coefficient described light-curve shape of magnitude (Stetson 1996);
24. `feature_weighted_mean_magn`, weighted mean magnitude (Malanchev et al. 2021);
25. `feature_anderson_darling_normal_flux`, unbiased Anderson–Darling normality test statistic for flux (Malanchev et al. 2021);
26. `feature_cusum_flux`, a range of cumulative sums dependent on the number of observations, mean flux, and flux standard deviation (Kim et al. 2014);
27. `feature_excess_variance_flux`, measure of the flux variability amplitude (Sánchez et al. 2017);
28. `feature_kurtosis_flux`, excess kurtosis of flux (Malanchev et al. 2021);
29. `feature_mean_variance_flux`, standard deviation of flux to mean flux ratio (Malanchev et al. 2021);
30. `feature_skew_flux`, skewness of flux (Malanchev et al. 2021);
31. `feature_stetson_k_flux`, Stetson K coefficient described light-curve shape of flux (Stetson 1996).

The full documentation, including equations, can be found here: https://docs.rs/light-curve-feature/0.2.2/light_curve_feature/features/index.html.

Appendix B
Host Galaxy Features

Our host galaxy features and a brief description are as follows:[76]

1. `gmomentXX`, second moment M_{xx} for g filter stack detection;
2. `gmomentXY`, second moment M_{xy} for g filter stack detection;
3. `gmomentYY`, second moment M_{yy} for g filter stack detection;
4. `gmomentR1`, first radial moment for g filter stack detection;
5. `gmomentRH`, half radial moment ($r^{0.5}$ weighting) for g filter stack detection;
6. `gPSFFlux`, PSF flux from g filter stack detection;
7. `gApFlux`, aperture flux from g filter stack detection;
8. `gKronFlux`, Kron (1980) flux from g filter stack detection;
9. `gKronRad`, Kron (1980) radius from g filter stack detection;
10. `gExtNSigma`, an extendedness measure for the g filter stack detection based on the deviation between PSF and Kron (1980) magnitudes, normalized by the PSF magnitude uncertainty;
11. `rmomentXX`, second moment M_{xx} for r filter stack detection;
12. `rmomentXY`, second moment M_{xy} for r filter stack detection;
13. `rmomentYY`, second moment M_{yy} for r filter stack detection;
14. `rmomentR1`, first radial moment for r filter stack detection;
15. `rmomentRH`, half radial moment ($r^{0.5}$ weighting) for r filter stack detection;
16. `rPSFFlux`, PSF flux from r filter stack detection;
17. `rApFlux`, aperture flux from r filter stack detection;
18. `rKronFlux`, Kron (1980) flux from r filter stack detection;
19. `rKronRad`, Kron (1980) radius from r filter stack detection;
20. `rExtNSigma`, an extendedness measure for the r filter stack detection based on the deviation between PSF and Kron (1980) magnitudes, normalized by the PSF magnitude uncertainty;
21. `imomentXX`, second moment M_{xx} for i filter stack detection;
22. `imomentXY`, second moment M_{xy} for i filter stack detection;
23. `imomentYY`, second moment M_{yy} for i filter stack detection;
24. `imomentR1`, first radial moment for i filter stack detection;
25. `imomentRH`, half radial moment ($r^{0.5}$ weighting) for i filter stack detection;
26. `iPSFFlux`, PSF flux from i filter stack detection;
27. `iApFlux`, aperture flux from i filter stack detection;
28. `iKronFlux`, Kron (1980) flux from i filter stack detection;
29. `iKronRad`, Kron (1980) radius from i filter stack detection;
30. `iExtNSigma`, an extendedness measure for the i filter stack detection based on the deviation between PSF and Kron (1980) magnitudes, normalized by the PSF magnitude uncertainty;
31. `zmomentXX`, second moment M_{xx} for z filter stack detection;
32. `zmomentXY`, second moment M_{xy} for z filter stack detection;
33. `zmomentYY`, second moment M_{yy} for z filter stack detection;
34. `zmomentR1`, first radial moment for z filter stack detection;

---

[76] Those from PS1 can be found here: https://outerspace.stsci.edu/display/PANSTARRS/PS1+Database+object+and+detection+tables.

35. `zmomentRH`, half radial moment ($r^{0.5}$ weighting) for z filter stack detection;
36. `zPSFFlux`, PSF flux from z filter stack detection;
37. `zApFlux`, aperture flux from z filter stack detection;
38. `zKronFlux`, Kron (1980) flux from z filter stack detection;
39. `zKronRad`, Kron (1980) radius from z filter stack detection;
40. `zExtNSigma`, an extendedness measure for the z filter stack detection based on the deviation between PSF and Kron (1980) magnitudes, normalized by the PSF magnitude uncertainty;
41. `ymomentXX`, second moment M_{xx} for y filter stack detection;
42. `ymomentXY`, second moment M_{xy} for y filter stack detection;
43. `ymomentYY`, second moment M_{yy} for y filter stack detection;
44. `ymomentR1`, first radial moment for y filter stack detection;
45. `ymomentRH`, half radial moment ($r^{0.5}$ weighting) for y filter stack detection;
46. `yPSFFlux`, PSF flux from y filter stack detection;
47. `yApFlux`, aperture flux from y filter stack detection;
48. `yKronFlux`, Kron (1980) flux from y filter stack detection;
49. `yKronRad`, Kron (1980) radius from y filter stack detection;
50. `yExtNSigma`, an extendedness measure for the y filter stack detection based on the deviation between PSF and Kron (1980) magnitudes, normalized by the PSF magnitude uncertainty;
51. `i-z`, aperture magnitude for the i filter stack detection minus aperture magnitude for the z filter stack detection; (iApMag − zApMag);
52. `gApMag_gKronMag`, aperture magnitude minus the Kron (1980) magnitude for the g filter stack detection;
53. `rApMag_rKronMag`, aperture magnitude minus the Kron (1980) magnitude for the r filter stack detection;
54. `iApMag_iKronMag`, aperture magnitude minus the Kron (1980) magnitude for the i filter stack detection;
55. `zApMag_zKronMag`, aperture magnitude minus the Kron (1980) magnitude for the z filter stack detection;
56. `yApMag_yKronMag`, aperture magnitude minus the Kron (1980) magnitude for the y filter stack detection;
57. `4DCD`, a four-dimensional color distance in $g-r$, $r-i$, $i-z$, and $z-y$ from the PS1 stellar locus, the path traced by stars in color–color space (Tonry et al. 2012);
58. `dist/DLR`, transient-host separation (arcseconds) normalized by the DLR.

Appendix C
Feature Correlations

With a large 120-dimensional feature-space, we are likely to have dependent or correlated features. In this subsection, we argue that this is an intentional choice that proffers quantitative benefits rather than suffers from poor design.

We present a Spearman rank-correlation matrix of our training set in Figure C1. Note that we choose Spearman's rank correlation in favor of Pearson correlation because the former can quantifiably describe linear and nonlinear relationships (those that can be described with a monotonic function), whereas the latter describes only linear relationships.

For light-curve features, we use the same set of 31 for both ZTF-r, ZTF-g bands. Broadly, the features that measure amplitude or amplitude-adjacent properties are correlated across intrapassband and interpassband. However, those that are correlated intrapassband are measuring different degrees of amplitude or amplitude variation during an SN's evolution (e.g., those measuring the 2nd, 10th, and 25th quantile of the magnitude distribution). Meanwhile, those that are interpassband correlated provide utility because of the color information they capture. Moreover, interpassband correlation vectors are less correlated than their intrapassband counterparts.

For host galaxy features, it is not only known that some features are highly correlated with each other; in fact, it has been leveraged in its use in empirical relationships such as the fundamental plane for ellipticals (Dressler et al. 1987), the color–magnitude relation (Bell et al. 2004), and more recently in cosmological analyses with SNe Ia in regard to host galaxy stellar mass (Popovic et al. 2021; Kelsey et al. 2023). See, e.g., Kelly et al. (2010), Sullivan et al. (2010), Brout & Scolnic (2021), and Grayling et al. (2024) for additional discussions regarding the relationship between SN and galaxy properties.

We find that, with the exception of the second moment `{g,r,i,z,y}momentXY`, all of our intrapassband host features are positively correlated, but to varying degrees. However, the same features interpassband tend to exhibit strong correlations (e.g., `{g,r,i,z,y}momentR1`). Easier to see is the appearance of blocks within the full correlation matrix, which are the strong positive correlations between brightness and radial moments of host galaxies in each band. These findings are largely consistent with those found in Gagliano et al. (2021; see further details in their Section 4).

From a broader viewpoint, the correlation between light-curve and host galaxy features is poor at best, with the largest positive Spearman correlation of $\rho \sim +0.30$ existing between the slope of the light curve in the least squares fit of the linear stochastic model `feature_linear_fit_slope_magn_g,r` and nearly all host galaxy features (with the exception of the second moment `{g,r,i,z,y}momentXY`, aperture magnitude $i-z$ color difference `i-z`, and normalized DLR `dist/DLR`). The most negative Spearman correlation of $\rho \sim -0.30$ relates the majority of host galaxy features to the nonweighted mean magnitude `feature_mean_magn_{g,r}` and the weighted mean magnitude `feature_weighted_mean_magn_{g,r}`. Perhaps the most interesting feature is the normalized DLR `dist/DLR`, which is not correlated with any other feature except for itself, and is the most discriminating host galaxy feature used in our RFC model (see Figure 6).

Typically, the large number of features used in this work could fall prey to the *curse of dimensionality* and would prohibit a brute-force ($\mathcal{O}(n)$ time) search across SNe. We circumvent this by (1) using an RFC for our AD model, which inherently performs feature selection when building each decision tree in the forest (where a random subset of `max_features` features at each split

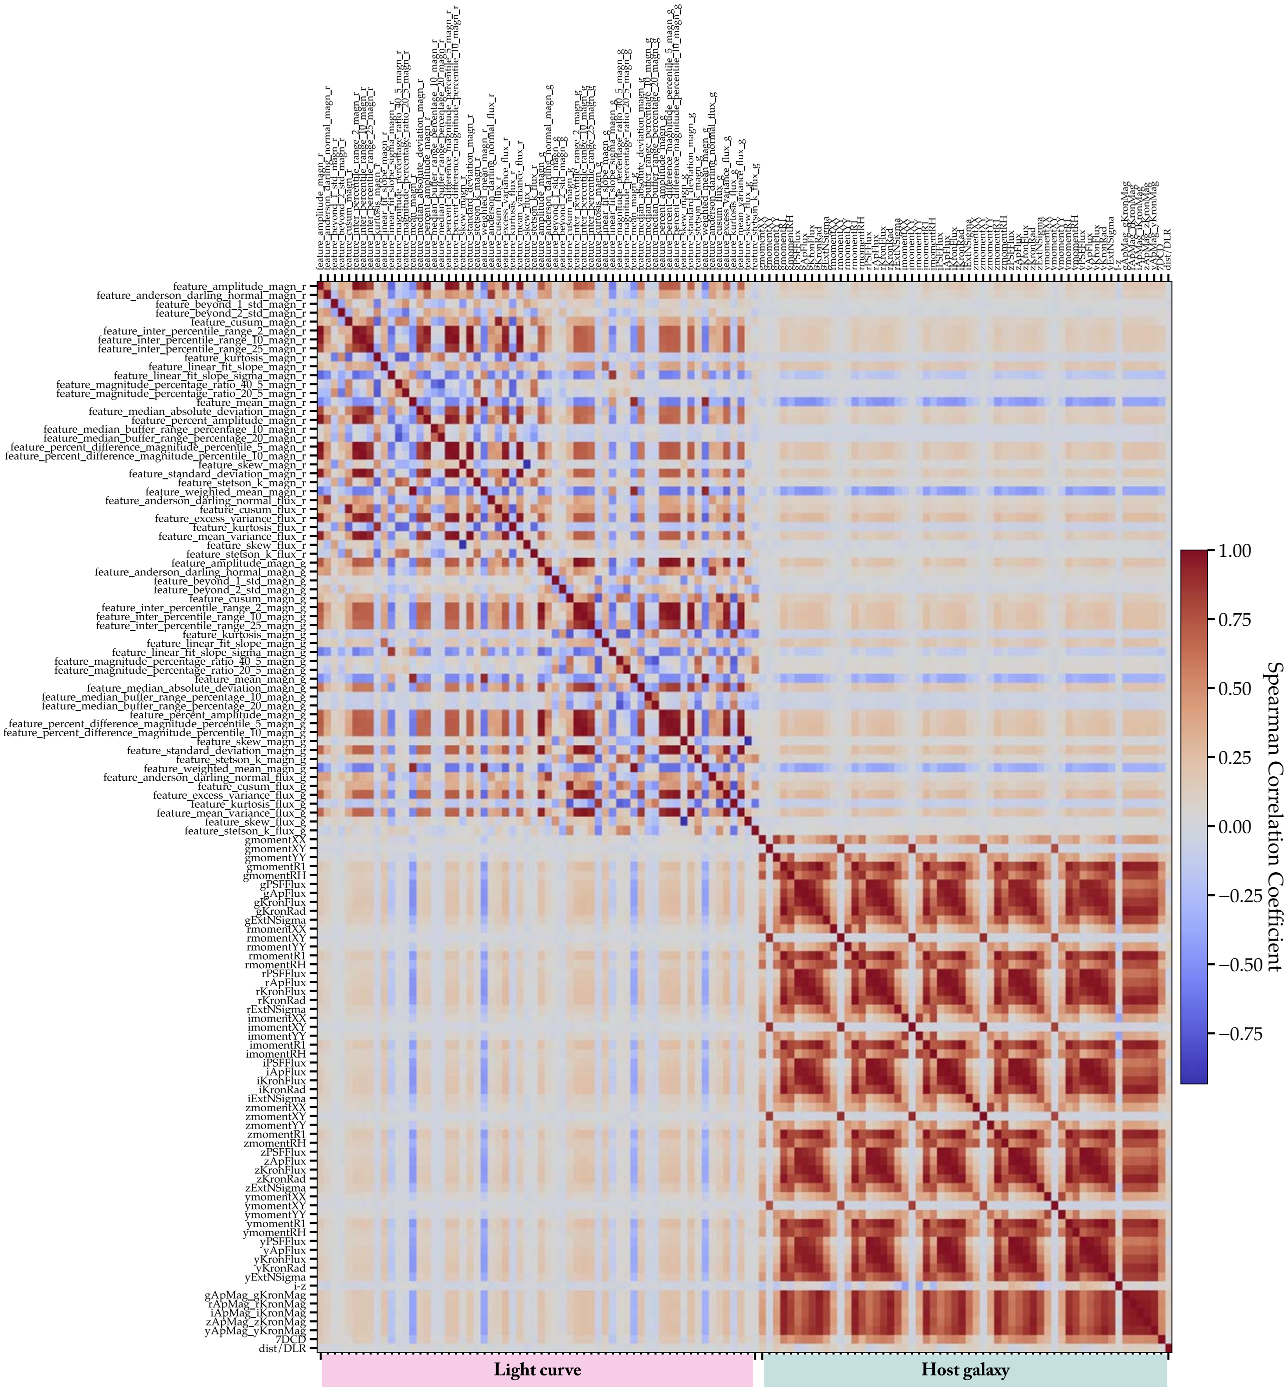

Figure C1. The Spearman rank-correlation matrix for the `LAISS` databank of extracted ZTF-*g*, ZTF-*r* light-curve and PS1-*grizy* host galaxy features (see Appendix B), with red corresponding to positively correlated features and blue corresponding to negatively correlated features. Broadly, the light-curve features that measure amplitude or amplitude-adjacent properties are correlated, more strongly intrapassband vs. interpassband. There exists strong positive correlations between brightness and radial moments of host galaxies in each band, forming a repeating block structure. From a high level, the correlation between light-curve features to host galaxy features is poor. Overall, these features are chosen to maximize color and capture related but different information of the supernova light curve and host galaxy, and to reduce the impact of dependent or irrelevant features.

is used[77]), and (2) using dimensionality reduction via PCA for our ANN similarity search with `ANNOY` (which itself is natively $\mathcal{O}(\log(n))$ time) to efficiently search the entire feature-space.

Appendix D
Methods for `PROSPECTOR` Host Galaxy SED Modeling

We correct photometry by subtracting an extinction correction using the corrected Schlegel, Finkbeiner, and Davis (SFD) map (Chiang 2023), provided through the `DUSTMAP`[78] package. To make these corrections, we assume an average $R_v = 3.1$. We set

---

[77] This process effectively reduces the dimensionality because not all features are used for each tree, helping to avoid overfitting, and is robust to more irrelevant features.

[78] https://github.com/gregreen/dustmaps

Table D1
`PROSPECTOR` Parameter Priors

Parameter	Prior Distribution	Notes
mass	$\log\mathcal{U}(1e6, 1e13)$	`PROSPECTOR` samples masses formed; we report stellar masses
logzsol	$\mathcal{U}(-2, 0.4)$	Limited by availability of MIST isochrones[a] (Choi et al. 2016)
dust2	$\mathcal{U}(0.0, 4.0)$	…
tage	$\mathcal{U}(0.0, 13.8)$	Age of stellar population
tau	$\log\mathcal{U}(0.01, 10)$	…
IMF	$\delta(1)$	Chabrier IMF
Dust Type	$\delta(1)$	Milky Way extinction law
SFH	$\delta(4)$	Delayed-Tau

Notes. We describe our model, initial parameter estimates, and priors below. We broadly follow the suggestions of model complexity following the procedure of Nugent et al. (2023). Below, $\mathcal{U}(a, b)$ is the top hat distribution on the interval (a, b). δ means that the parameter is always fixed.
[a] The MIST models developed for ancient, metal-poor populations from $-4.0 \leqslant [Z/H] < -2.0$ from Choi et al. (2016) are not included in `FSPS`.

a lower bound on the photometric error in any filter as 15% to allow for systematic error—which we note is a larger bound than used in Leja et al. (2019b). We felt it was necessary to increase the lower bound of uncertainty given our inexact matched aperture sizes.

The final vector of photometric observations and uncertainties are used to fit our stellar population synthesis model. Prospector calls the `PYTHON-FSPS`[79] library to perform the fits (Conroy et al. 2009; Conroy & Gunn 2010). The library uses the MIST isochrones to build the models (Choi et al. 2016; Dotter 2016). We use a Chabrier initial mass function (IMF; Chabrier 2003) with a delayed-tau stellar formation history ($t \times \exp{-t/\tau}$). We set the redshift to the value of the known transient's redshift. We model dust attenuation using the Milky Way dust extinction law (Cardelli et al. 1989), and additionally attenuate young stars with an extra static dust contribution of 0.5 (for additional details, see Conroy et al. 2009). We model the contribution from nebular emission but fix the gas phase metallicity and the gas ionization parameters (Byler et al. 2017). We use the Gallazzi et al. (2005) mass–metallicity relationship as a prior on the stellar metallicity, following its use in Nugent et al. (2023). The final model is a function of five free parameters. Our priors are described in Table D1.

Appendix E
Spectra

Here, we describe our observations and data reduction for the nine spectra whose classifications are used in Table 5.

TDE 2023kvy was discovered on 2023 June 16 23.59 by Tonry et al. (2023) and classified as a TDE on 2023 September 9.97 (Gomez et al. 2023b) as part of the FLEET program (Gomez et al. 2020) from a SOAR/Goodman spectrum acquired on 2023 August 27. `LAISS` correctly tagged TDE 2023kvy as an anomaly on 2023 August 26. This spectrum was acquired in coordination with FLEET, whose program identified this object independently.

SN 2023sds was discovered on 2023 September 11.38 by Munoz-Arancibia et al. (2023b) and classified by us as an SN II on 2024 March 24.10 (Aleo 2024) from a Lick 3 m/KAST spectrum acquired by us on 2023 October 22.19. `LAISS` correctly tagged SN 2023sds as an anomaly on 2023 September 24.

SN 2023nof was discovered on 2023 July 23.34 by Munoz-Arancibia et al. (2023c) and originally classified as an SN II on 2023 July 27.34 by Lidman et al. (2023) from an Australian National University (ANU) 2.3 m/WiFeS spectrum acquired on 2023 July 26. Later, we reclassified as an SN IIn on 2023 October 6.77 from an ANU 2.3 m/WiFeS spectrum acquired by us on 2023 September 29 (Aleo et al. 2023). `LAISS` correctly tagged SN 2023nof as an anomaly on 2023 September 24.

SN 2023nwe was discovered on 2023 July 24.43 by Chambers et al. (2023) and classified by us as an SN IIn on 2023 September 11.69 (Davis et al. 2023) from a Lick 3 m/KAST spectrum acquired by us on 2023 September 6.30. `LAISS` correctly tagged SN 2023nwe as an anomaly on 2023 August 28.

SN 2023omf was discovered on 2023 July 31.35 by Fremling (2023a) and classified by us as an SN IIn on 2024 March 24.15 (Aleo 2024) from a Lick 3 m/KAST spectrum acquired by us on 2023 October 22.11. `LAISS` correctly tagged SN 2023nwe as an anomaly on 2023 October 10.

SN 2023ocx was discovered on 2023 July 20.22 by Fremling (2023b) and classified as an SN Ia-CSM by us on 2023 October 2.52 (Angus & Aleo 2023) from an NOT/ALFOSC spectrum acquired by us on 2023 September 19. `LAISS` correctly tagged SN 2023ocx as an anomaly on 2023 August 28.

SN 2023otw was discovered on 2023 July 29.46 by Fremling (2023c) and classified as an SN II on 2023 November 7.56 (Gomez et al. 2023c) as part of the FLEET program (Gomez et al. 2020) from a SOAR/Goodman spectrum acquired on 2023 October 11. This spectrum was acquired in coordination with FLEET, whose program identified this object independently. `LAISS` incorrectly tagged SN 2023otw as an anomaly on 2023 September 20.

SN 2023swf was discovered on 2023 September 16.25 by Fremling (2023d) and classified as an SN II on 2023 October 4.72 (Ayala et al. 2023) from a ESO-New Technology Telescope (NTT) /EFOSC2-NTT (ePESSTO+) spectrum acquired on 2023 October 04.21. We retained the original SN II classification from an ANU 2.3 m/WiFeS spectrum acquired by us on 2023 October 11. `LAISS` incorrectly tagged SN 2023swf as an anomaly on 2023 October 9.

SN 2023sed was discovered on 2023 September 8.35 by Munoz-Arancibia et al. (2023b) and classified by us as an SN Ia on 2023 October 2.55 (Angus 2023) from an NOT/ALFOSC spectrum acquired by us on 2023 September 28. `LAISS` incorrectly tagged SN 2023sed as an anomaly on 2023 September 24.

To reduce the Kast data, we used the `UCSC Spectral Pipeline`[80] (Siebert et al. 2020), a custom data-reduction pipeline based on procedures outlined by Foley et al. (2003), Silverman et al. (2012), and references therein. The two-dimensional spectra were bias-corrected, flat-field corrected,

---

[79] https://github.com/dfm/python-fsps

[80] https://github.com/msiebert1/UCSC_spectral_pipeline

Table E1
Spectroscopic Observation Details

IAU Name	Obs Date (UT)	MJD	Estimated Phase (days)	Telescope	Instrument
2023kvy	2023-08-27	60183	+65	SOAR	GOODMAN
2023sds	2023-10-22	60239	+34	Lick 3 m	KAST
2023nof	2023-09-29	60196	+43	ANU 2.3 m	WiFeS
2023nwe	2023-09-06	60193	+14	Lick 3 m	KAST
2023omf	2023-10-22	60239	+52	Lick 3 m	KAST
2023ocx	2023-09-19	60206	+33	NOT	ALFOSC
2023otw	2023-10-11	60228	+65	SOAR	GOODMAN
2023swf	2023-10-11	60228	+7	ANU 2.3 m	WiFeS
2023sed	2023-09-28	60215	+4	NOT	ALFOSC

Note. Log of spectroscopic observations for TDE 2023kvy, SN 2023sds, SN 2023nof, SN 2023nwe, SN 2023omf, SN 2023ocx, SN 2023otw, SN 2023swf, SN 2023sed.

adjusted for varying gains across different chips and amplifiers, and trimmed. One-dimensional spectra were extracted using the optimal algorithm (Horne 1986). The spectra were wavelength calibrated using internal comparison-lamp spectra with linear shifts applied by cross-correlating the observed night-sky lines in each spectrum to a master night-sky spectrum. Flux calibration and telluric correction were performed using standard stars at a similar airmass to that of the science exposures. We combine the sides by scaling one spectrum to match the flux of the other in the overlap region and use their error spectra to correctly weight the spectra when combining. More details of this process are discussed elsewhere (Foley et al. 2003; Silverman et al. 2012; Siebert et al. 2020; Davis et al. 2023).

Data obtained with ALFOSC, GOODMAN, and WiFeS were reduced using standard techniques, which included correction for bias, overscan, and flat-field. Spectra of comparison lamps and standard stars acquired during the same night and with the same instrumental setting have been used for the wavelength and flux calibrations, respectively. We employed standard IRAF commands to extract all spectra.

A table summarizing all nine spectroscopic follow-up observations follows in Table E1.

Appendix F
Tables

In Table F1, we report the following previously undiscovered SNe to TNS from ANN = 8 nearest neighbor matches to all SNe classes. Bold text indicates an object we consider a likely anomaly. Italicized text indicates an object that appeared more than once from the ANN = 8 search (e.g., ZTF18abmofew/2018mfo was uniquely found seven times during the ANN = 8 search). These objects are listed for completeness for each match shown. Only classes for which at least one discovery was made are shown.

In Table F2, we report the following previously undiscovered SNe to TNS from an exhaustive search of our databank.

Table F1
84 Previously Undiscovered/Reported SNe Found by ANN = 8 Search

ZTF ID	IAU Name	Possible Class	Reference SNe	ANN	Remarks
Reference Class: SN Ia (1098 objects)					
ZTF18achdfqm	2018mfl	SN Ia?	2020vnl	4	Host-$z = 0.095$. Peak $M_{\rm abs} \sim -18.8$ mag.
ZTF18abnzney	*2018mfm*	SN Ia?	2020whs	1	⋯
ZTF21abcluco	*2021ahvn*	SN Ia?	”	3	⋯
ZTF17aadqidc	2021ahvo	SN I?	”	5	Host-$z = 0.087$. Peak $M_{\rm abs} \sim -18.6$ mag.
ZTF18acpegrg	2018mfn	SN Ia?	2020wts	2	Host-$z = 0.132$. Peak $M_{\rm abs} \sim -19.6$ mag.
ZTF18abmofew	*2018mfo*	SN Ia?	2020xyh	2	Well sampled.
ZTF19aaaamwp	2018mfp	SN Ia?	2020xit	5	Host-$z = 0.124$. Peak $M_{\rm abs} \sim -19.1$ mag.
ZTF21aamssts	*2021ahvp*	SN Ia?	2020zcw	5	⋯
ZTF21abvjcnb	2021ahvq	SN I?	”	7	⋯
ZTF18acbvtjm	*2018mfq*	SN?	2020zbr	1	⋯
ZTF20aconebc	2020afeq	SN?	”	5	⋯
ZTF21aamssts	*2021ahvp*	SN Ia?	2020abrg	1	⋯
ZTF18abrxvpd	2018mfr	SN Ia?	”	4	⋯
ZTF20adadbsm	*2021ahvr*	SN I?	2020acmi	5	⋯
ZTF20adadbsm	*2021ahvr*	SN I?	2020acvz	4	⋯
ZTF21aalgilf	2021ahvs	SN Ia?	2021ab	8	⋯
ZTF21aarpnxt	2021ahvt	SN Ia?	2021Y	1	⋯
ZTF18aczerlj	*2018mfs*	SN II/IIP?	2021by	2	Well sampled. Host-$z = 0.066$. Peak $M_{\rm abs} \sim -17.7$ mag.
ZTF18acbzvzm	*2018mft*	SN Ia?	2021vt	1	Host-$z = 0.106$. Peak $M_{\rm abs} \sim -19.5$ mag.
ZTF21aamssts	*2021ahvp*	SN Ia?	2021apk	8	⋯
ZTF18acbwdym	*2018mfu*	SN Ia?	2021arr	2	Host-$z = 0.091$. Peak $M_{\rm abs} > -19.0$ mag.
ZTF18acbwgmi	2018mfv	SN Ib/c?	2021buy	5	Tagged by AD model.
ZTF19abxsehw	2019aata	SN I?	2021eij	3	⋯
ZTF20acvdqsy	2020afer	SN Ia?	”	7	⋯
ZTF18abokvkt	*2018mfw*	SN Ia?	2021gfi	3	Faint host.
ZTF19aaviczu	2019aatb	SN Ia?	2021iok	2	Host-$z = 0.156$. $M_{\rm abs} \sim -19.6$ mag.
ZTF18abtgmuw	2018mfx	SN Ia?	”	8	Faint host.
ZTF18abnvnqb	2018mfy	SN Ia?	2021lea	4	Well sampled. Faint host.
ZTF21aalimvt	2021ahvu	SN?	2021mhm	2	No decline.
ZTF21aadruss	2021ahvv	SN Ia?	2021mla	4	⋯
ZTF21aaxswjy	2021ahvw	SN Ia?	2021nya	5	Faint host.
ZTF21aajtfas	*2021ahvx*	SN Ia?	2021oia	1	Well sampled.
ZTF21aaswuea	2021ahvy	SN Ia?	2021nsh	7	Host-$z = 0.134$. $M_{\rm abs} \sim -19.0$ mag.
ZTF21aajtfas	*2021ahvx*	SN Ia?	2021tvc	5	Well sampled.
ZTF18abmofew	*2018mfo*	SN Ia?	2021uwa	5	Well sampled.
ZTF20acuyxki	2020afes	SN Ia?	2021uvx	1	Host-$z = 0.121$. $M_{\rm abs} \sim -19.2$ mag.
ZTF18abmofew	*2018mfo*	SN Ia?	2021vku	4	Well sampled.
ZTF18abtjrbt	2018mfz	TDE?	2021vnv	7	Blue. SALT3 $x_1 = +3.00$. Host-$z = 0.065$. $M_{\rm abs} \sim -18.4$ mag.
ZTF18ablqjws	*2018mga*	SN II/IIn?	2021wym	3	Well sampled. Tagged by AD model.
ZTF20acjoazg	*2020afet*	SN Ia?	2021xoh	6	Well sampled.
ZTF20acthdtq	2020afeu	SN Ia?	2021wwc	1	Faint host.
ZTF18acbzvzm	*2018mft*	SN Ia?	2021xpk	1	Host-$z = 0.106$. Peak $M_{\rm abs} \sim -19.5$ mag.
ZTF18abmofew	*2018mfo*	SN Ia?	2021ycy	5	Well sampled.
ZTF18abuioue	*2018mgb*	SN Ia?	”	2	No visible host.
ZTF20acjoazg	*2020afet*	SN Ia?	2021ysn	2	Well sampled.
ZTF18acvgwfq	2018mgc	SN Ia?	2021aafd	1	Fast decline.
ZTF18acbwdym	*2018mfu*	SN Ia?	2021aaem	2	Host-$z = 0.091$. Peak $M_{\rm abs} > -19.0$ mag.
ZTF18abwhsnx	*2018mgd*	SN Ib/c?	2021aaht	7	Tagged by AD model.
ZTF18abwhsnx	*2018mgd*	SN Ib/c?	2021abds	2	Tagged by AD model.
ZTF18abmofew	*2018mfo*	SN Ia?	2021abaq	7	Well sampled.
ZTF21abdhwhj	2021ahvz	SN?	2021abpc	4	⋯
ZTF21aamssts	*2021ahvp*	SN Ia?	2021acgo	5	⋯
ZTF18acpefbb	2018mge	SN Ia?	2017baq	4	Host-$z = 0.135$. Peak $M_{\rm abs} \sim -19.5$ mag.
ZTF21aaualdi	2021ahwa	SLSN?	2021acgu	7	No visible host. FLEET = 46%(41%) SLSN-I(II). Tagged by AD model.
ZTF18absljwl	*2018mgf*	SN II?	2021acfg	6	No rise.
ZTF18acbzvzm	*2018mft*	SN Ia?	2021acwr	1	Host-$z = 0.106$. Peak $M_{\rm abs} \sim -19.5$ mag.
ZTF18acefgee	2018mgg	SN I?	2021aczd	1	⋯
ZTF21aabvpub	*2021ahwb*	SN Ia?	”	7	⋯
ZTF21aayngti	2021ahwc	SN Ia?	”	8	Host-$z = 0.142$. Peak $M_{\rm abs} \sim -19.1$ mag.
ZTF21abcluco	*2021ahvn*	SN Ia?	2021adeq	2	⋯
ZTF18acbvijf	*2018mgh*	SN Ia?	”	5	⋯
ZTF18abuioue	*2018mgb*	SN Ia?	2021aden	6	No visible host.

Table F1
(Continued)

ZTF ID	IAU Name	Possible Class	Reference SNe	ANN	Remarks
ZTF19aacislx	2019aatc	SN Ia?	2021aefa	1	Faint host.
ZTF21aamfdzo	*2021ahwd*	SN Ia?	2022is	8	⋯
ZTF21abghbue	2021ahwe	SN Ia?	2022aby	4	Faint host.
ZTF18acwward	2018mgi	SN Ia?	2022bcc	7	Faint host.
ZTF20acpgnmw	2020afev	TDE?	2022ccz	7	Blue. Nuclear. FLEET = 63% TDE.
ZTF18acbvtjm	*2018mfq*	SN?	2022cmd	3	⋯
ZTF21aabvpub	*2021ahwb*	SN Ia?	2022csq	3	⋯
ZTF18abuluko	2018mgj	SN Ia?	”	4	No decline.
Reference Class: SN II (243 objects)					
ZTF18ablqjws	*2018mga*	SN II/IIn?	2021lmp	7	Well sampled. Tagged by AD model.
ZTF20acyybza	2021ahwf	SN II?	2018grp	5	Well sampled. Faint host.
ZTF20acbiwfi	*2020afew*	SLSN?	2019cob	4	FLEET = 44% SLSN-II.
ZTF20aczhbki	2020afex	SLSN?	”	5	Host-$z = 0.573$. Peak $M_{\rm abs} \sim -23.1$ mag.
ZTF20acyzxse	2020afey	SN Ia?	”	7	⋯
ZTF18abkhqww	2018mgk	SN II/IIP?	2020tab	2	Well sampled. Faint host.
ZTF18acandwv	2018mgl	SN Ia?	2020tzs	1	⋯
ZTF18aczerlj	*2018mfs*	SN II/IIP?	2020vcp	5	Well sampled. Host-$z = 0.066$. Peak $M_{\rm abs} \sim -17.7$ mag.
ZTF21aaiapis	2021ahwg	SN Ib/c?	2021adcw	4	Tagged by AD model.
ZTF18aaxckpt	2018mgm	SN?	2020yae	7	⋯
ZTF18aczerlj	*2018mfs*	SN II/IIP?	2020acjg	3	Well sampled. Host-$z = 0.066$. Peak $M_{\rm abs} \sim -17.7$ mag.
ZTF20acyydhn	2020afez	SN II?	2021crx	3	Host-$z = 0.021$. Peak $M_{\rm abs} \sim -15.3$ mag.
ZTF21abasbpe	2021ahwh	SN IIn/SLSN?	2021uuz	4	Faint host. No decline. FLEET = 61% SLSN-II.
ZTF18abmofew	*2018mfo*	SN Ia?	2021yky	3	⋯
ZTF21aajtfas	*2021ahvx*	SN Ia?	2021aaev	8	Well sampled.
Reference Class: SN IIn (59 objects)					
ZTF18ablqjws	*2018mga*	SN II/IIn?	2021lft	1	Well sampled. Tagged by AD model.
ZTF19aasalud	2019aatd	SN Ia?	2021acnp	6	⋯
Reference Class: SN Ia-91T-like (37 objects)					
ZTF21aalvdng	2021ahwi	SN Ia?	2022cvt	6	No decline.
ZTF21aamfdzo	*2021ahwd*	SN Ia?	2020veg	7	⋯
ZTF18abokvkt	*2018mfw*	SN Ia?	2020wze	8	Faint host.
ZTF20aclghmy	*2020affa*	SN IIn?	2021hj	5	SALT3 $x_1 = +3.00$. Host-$z = 0.136$. Peak $M_{\rm abs} \sim -19.4$ mag.
Reference Class: SN Ib (25 objects)					
ZTF18acbvijf	*2018mgh*	SN Ia?	2021lax	5	⋯
ZTF18abwhsnx	*2018mgd*	SN Ib/c?	2021riw	6	Tagged by AD model.
ZTF18acrdqvy	2018mgn	SN II?	2021aghp	3	Host-$z = 0.049$. Peak $M_{\rm abs} \sim -17.0$ mag.
ZTF19aabblsx	2019aate	SN Ia?	”	7	⋯
Reference Class: SN IIb (21 objects)					
ZTF18aaiuynw	2018mgo	SN Ia?	2018hqu	3	⋯
ZTF18acenyrb	2018mgp	SN I?	2020abkp	4	⋯
Reference Class: SN Ic (21 objects)					
ZTF21abvatnb	2021ahwj	SN Ia?	2021lei	3	⋯
ZTF21aagyuvz	2021ahwk	SN I?	”	8	⋯
ZTF19aafmxxd	2019aatf	SN Ia?	2021acwh	2	Host-$z = 0.095$. Peak $M_{\rm abs} \sim -19.1$ mag.
ZTF21aajtfas	*2021ahvx*	SN Ia?	2021adgu	2	Well sampled.
ZTF ID	IAU Name	Possible class.	Reference SNe	ANN	Remarks
Reference Class: TDE (20 objects)					
ZTF20acqyjih	2020affb	SN Ia?	2021lo	7	Nuclear.
ZTF21aasdcgt	2021ahwl	TDE?	2021gje	6	FLEET = 60% TDE. Nuclear, blue. Tagged by AD model.
ZTF21acbgyai	2021ahwm	SN Ia?	2022rz	2	Nuclear.
Reference Class: SN IIP (14 objects)					
ZTF18absljwl	*2018mgf*	SN II?	2020tet	7	No rise.
Reference Class: SLSN-II (14 objects)					
ZTF18abuioue	*2018mgb*	SN Ia?	2020vws	7	No visible host.
ZTF21aajtfas	*2021ahvx*	SN Ia?	2021fmu	4	Well sampled.
ZTF20acllkua	2020affc	SN Ia?	2022akb	1	⋯
ZTF18aajmlnp	2019aatg	SN Ia?	”	1	Host-$z = 0.117$. Peak $M_{\rm abs} \sim -19.1$ mag.

Table F1
(Continued)

ZTF ID	IAU Name	Possible Class	Reference SNe	ANN	Remarks
			Reference Class: SN Ic-BL (14 objects)		
ZTF21aalimtb	2021ahwn	SN I?	2021otc	3	⋯.
			Reference Class: SLSN-I (11 objects)		
ZTF18abuioue	*2018mgb*	SN Ia?	2020xga	4	No visible host.
ZTF20acteioa	2020affd	SN Ia?	2021rwz	3	Faint host.
ZTF18abwfhqb	2018mgq	SN II?	2021ybf	5	Faint host.
ZTF19aaapmth	2018mgr	SN Ia?	”	6	Faint host.
ZTF18abnzney	*2018mfm*	SN Ia?	2021ybf	7	⋯
			Reference Class: SN Ia-pec (10 objects)		
ZTF19aaptarf	2019aath	SN Ia?	2021wwu	5	⋯
ZTF18acvgehz	2018mgs	SN Ia?	2022bbt	6	Host-$z = 0.120$. Peak $M_{\rm abs} \sim -19.4$ mag.
			Reference Class: AGN (seven objects)		
ZTF18abmofew	*2018mfo*	SN Ia?	2021swi	4	Well sampled.
			Reference Class: SN Ibn (seven objects)		
ZTF18acbwdym	*2018mfu*	SN Ia?	2021bbv	2	Host-$z = 0.091$. Peak $M_{\rm abs} > -19.0$ mag.
ZTF20acbiwfi	*2020afew*	SLSN?	”	7	FLEET = 44% SLSN-II.
			Reference Class: SN (five objects)		
ZTF18acueeoo	2018mgt	SN II?	2020uts	6	Host-$z = 0.085$. Peak $M_{\rm abs} \sim -18.2$ mag. Tagged by AD model.
			Reference Class: SN IIn-pec (one object)		
ZTF20aclghmy	*2020affa*	SN IIn?	2021vlu	5	SALT3 $x_1 = +3.00$. Host-$z = 0.136$. Peak $M_{\rm abs} \sim -19.4$ mag.

Notes. Bold text indicates an object we consider a likely anomaly. Italicized text indicates an object that appeared more than once from the ANN = 8 search. This is explained in the text preceding the table. Parentheses indicate the same object in the row above for that column.

Table F2
241 Previously Undiscovered SNe Found by Exhaustive Search

ZTF ID	IAU Name	Possible Class	Remarks
ZTF18aabteyx	2018mgu	SN?	SN in template thumbnail.
ZTF18aajgowk	2018mgw	SN Ib/c?	Tagged by AD model.
ZTF18aajkgtr	2018mgx	SN Ia?	…
ZTF18aamdjfh	2018mgv	SN?	SN in template thumbnail.
ZTF18aaqzdge	2018mgy	SN?	No rise.
ZTF18aathofv	2018mgz	SN Ia?	No rise.
ZTF18aawlezq[a]	2018mha	SN II?	…
ZTF18aawonvg	2018mhb	SN Ia?	…
ZTF18aaxmjit	2018mhc	SN?	No rise.
ZTF18aaxqwnd	2018mhd	SN II?	…
ZTF18aaxzhzf	2018mhe	SN Ia?	…
ZTF18aayyedm	2018mhf	SN Ia?	…
ZTF18abhttmr	2018mhg	SN?	No visible host.
ZTF18abiitmq	2018mhh	SN II?	Long ($\sim$70 day) rise. Candidate member of long-rising SN II class. No decline.
ZTF18abjicev	2018mhi	SN II?	Lasts $\sim$200 days.
ZTF18abkmfaj	2018mhj	SN II?	Long ($\sim$50 day) rise. Candidate member of long-rising SN II class. No decline.
ZTF18ablwatx	2018mhk	SN I?	Edge-on host galaxy.
ZTF18abnysyy	2018mhl	SN Ia?	Well sampled.
ZTF18abolrnr	2018mhm	SN Ia?	…
ZTF18abpucqe	2018mhn	SN?	…
ZTF18abrlhnm	2018mho	SN Ia?	Faint host.
ZTF18abrlurr	2018mhp	SN I?	…
ZTF18abrwqti	2018mhq	SN Ia?	…
ZTF18abscetq	2018mhr	SN Ia?	…
ZTF18abscghc	2018mhs	SN Ia?	…
ZTF18absdgly	2018mht	SN I?	Faint host.
ZTF18abshfmb	2018mhu	SN II?	Faint host.
ZTF18abskrcm	2018mhv	SN Ia?	…
ZTF18absljlb	2018mhw	SN II/IIP?	Faint host.
ZTF18absllcn	2018mhx	SN Ia?	…
ZTF18absloog	2018mhy	SN Ia?	…
ZTF18absmsbm	2018mhz	SN Ia?	Faint host.
ZTF18absopva	2018mia	SN?	…
ZTF18absqohp	2018mib	SN Ia?	No visible host.
ZTF18absquza	2018mic	SN I?	Edge-on host galaxy.
ZTF18absrljp	2018mid	SN Ia?	…
ZTF18abtflop	2018mie	SN Ia?	…
ZTF18abtigie	2018mif	SN?	…
ZTF18abtlxae	2018mig	SN Ia?	…
ZTF18abtmnha	2018mih	SN Ia?	…
ZTF18abtnwpa	2018mii	SN Ia?	…
ZTF18abtotsq	2018mij	SN Ia?	No visible host.
ZTF18abtpgms	2018mik	SN?	…
ZTF18abtptey	2018 mil	SN?	…
ZTF18abtqceg	2018mim	SN Ia?	Faint host.
ZTF18abtqidt	2018 min	SN Ia?	…
ZTF18abuahio	2018mio	SN Ia?	No visible host. Tagged by AD model.
ZTF18abubjqi	2018mip	SN Ia?	…
ZTF18abufwxs	2018miq	SN I?	Faint host. No decline.
ZTF18abuvqgo	2018mir	SN Ia?	…
ZTF18abuxgmk	2018mis	SN?	…
ZTF18abuyomg	2018mit	SN Ia?	…
ZTF18abvctat	2018miu	SN?	No visible host.
ZTF18abvexiz	2018miv	SN Ia?	Faint host.
ZTF18abvrgjc	2018miw	SN Ia?	Faint host.
ZTF18abvywcl	2018mix	SN Ia?	Faint host.
ZTF18abwblyl	2018miy	SN Ia?	…
ZTF18abwqgsc	2018miz	SN Ia?	…
ZTF18abwsfdy	2018mja	SN II/IIP?	Faint host.
ZTF18abxfott	2018mjb	SN Ia?	Faint host.
ZTF18abxthsc	2019aati	SN I?	…
ZTF18acajloc	2018mjc	SN?	Faint host. Tagged by AD model.
ZTF18acbvgtj	2018mjd	SN?	…
ZTF18acbvzsi	2018mje	SN Ia?	Faint host.
ZTF18acbwele	2018mjf	SN?	No rise.
ZTF18acbxrft	2018mjg	SN Ia?	…
ZTF18acbzvfj	2018mjh	SN?	Faint host.
ZTF18acbzvow	2018mji	SN?	…
ZTF18acchzlc	2018mjj	SN?	Faint host.
ZTF18accjwbc	2018mjk	SN?	Edge-on host galaxy.
ZTF18accjyim	2018mjl	SN?	Faint host.
ZTF18acckcfr	2018mjm	SN?	…
ZTF18accndre	2018mjn	SN?	Faint host.
ZTF18accvkpt	2018mjo	SN Ia?	…
ZTF18accwild	2018mjp	SN?	Faint host.
ZTF18acegbeo	2018mjq	SN Ia?	Edge-on host galaxy.
ZTF18acehtvv	2018mjr	SN?	…
ZTF18aceitqw	2018mjs	SN?	Faint host.
ZTF18acevgyz	2018mjt	SN I?	Tagged by AD model.
ZTF18achdign	2018mju	SN?	…
ZTF18acmyfyu	2018mjv	SN Ia?	…
ZTF18acmymdx	2018mjw	SN?	…
ZTF18acnmicp	2018mjx	SN?	…
ZTF18acnneyt	2018mjy	SN II/IIP?	Faint host.
ZTF18acnnfqr	2018mjz	SN Ia?	Host-$z = 0.171$. Peak $M_{\rm abs} \sim -19.4$ mag.
ZTF18acqzakw	2018mka	SN Ia?	Faint host.
ZTF18acurkik	2018mkb	CC SN?	SALT3 $x1 = +3.00$. Host-$z = 0.068$. Peak $M_{\rm abs} \sim -17.9$ mag.
ZTF18acvgmpx	2018mkc	SN II?	Host-$z = 0.069$. Peak $M_{\rm abs} \sim -17.6$ mag. In ring galaxy?
ZTF18acvwkvc	2018mkd	TDE?	FLEET = 40% TDE, 27% SLSN-II.
ZTF18acwtrfe	2018mke	SN?	…
ZTF18acwworr	2018mkf	SN II?	Host-$z = 0.092$. Peak $M_{\rm abs} \sim -18.3$ mag.
ZTF18acwyauj	2018mkg	SN?	…
ZTF18acyxxen	2018mkh	SN Ia?	…
ZTF18aczzvtw	2018mki	SN?	Faint host.
ZTF18aczzwjf	2018mkj	SN?	Faint host.
ZTF18adaksjk	2018mkk	SN?	Faint host.
ZTF18adazgwk	2018mkl	SN?	…
ZTF18adcbkiq	2018mkm	SN?	…
ZTF19aaaefkt	2018mkn	SN Ia?	…
ZTF19aaaeuxz	2019aatj	SN I?	…
ZTF19aaafmjk	2019aatk	SN II?	Host-$z = 0.062$. Peak $M_{\rm abs} \sim -17.6$ mag.
ZTF19aaafotc	2019aatl	SN Ia?	Host-$z = 0.137$. Peak $M_{\rm abs} \sim -19.6$ mag.
ZTF19aaajbre	2018mko	SN Ia?	…
ZTF19aaapdum	2019aatm	SN Ia?	Host-$z = 0.171$. Peak $M_{\rm abs} > -19.6$ mag.

Table F2
(Continued)

ZTF ID	IAU Name	Possible Class	Remarks
ZTF19aacwiqw	2019aatn	SN Ia?	⋯
ZTF19aacyrdk	2019aato	SN Ia?	Faint host.
ZTF19aadgceq	2019aatp	SN Ia?	⋯
ZTF19aavoqbe	2019aatq	SN Ia-91bg-like?	Host-$z = 0.13$. SALT3 $c = +0.3$. Peak $M_{\rm abs} \sim -18.8$ mag.
ZTF19abiinmg	2019aatr	SN Ia?	⋯
ZTF19abljudj	2019aats	SN II?	In elliptical galaxy?
ZTF20abribtl	2020affe	SN?	⋯
ZTF20abztvjb	2020afff	SN II/IIP?	Red. Lasted $\sim$80 days. Faint host.
ZTF20acbptqx	2020affg	SN IIn?	Faint host. FLEET = 52% SN II, 31% SLSN-II.
ZTF20acfamrq	2020affh	SN II?	⋯
ZTF20achofuu	2021ahwo	SN Ia?	⋯
ZTF20acipeqz	2020affi	SN II?	⋯
ZTF20ackjgbt	2020affj	SN I?	Faint host.
ZTF20acmsesb	2020affk	SN II?	⋯
ZTF20acnbhsq	2020affl	SN Ia?	⋯
ZTF20acnexgl	2020affm	SN Ia?	Host-$z = 0.118$. Peak $M_{\rm abs} \sim -19.2$ mag.
ZTF20acngnvf	2020affn	SN IIn?	Faint host. Lasted 125 days. FLEET = 70% SN II.
ZTF20acnviuc	2020affo	SN Ia?	⋯
ZTF20acnvoql	2020affp	SN Ia?	⋯
ZTF20acnzpsj	2020affq	SN Ia?	⋯
ZTF20acnztaa	2020affr	SN II/IIn?	Faint host.
ZTF20acnzzlo	2020affs	SN Ia?	⋯
ZTF20acotzgs	2020afft	SN Ia?	Faint host.
ZTF20acpggpe	2020affu	SN Ia?	⋯
ZTF20acpskkx	2020affv	SN?	⋯
ZTF20acpxtkj	2020affw	SN Ia?	⋯
ZTF20acpzjsk	2020affx	TDE?	Blue. Nuclear. FLEET = 40% SN II.
ZTF20acrktth	2020affy	SN Ia?	⋯
ZTF20acrssoi	2020affz	SN IIn?	Faint host. FLEET = 74% SN II.
ZTF20acsfdpi	2020afga	SN?	⋯
ZTF20actkqax	2020afgb	SN IIn/SLSN?	Faint host. FLEET = 40% SN II, 36% SLSN-II. Tagged by AD model.
ZTF20acuhlsk	2020afgc	SN Ia?	⋯
ZTF20acusjnl	2020afgd	SN Ia?	Faint host.
ZTF20acxyrkk	2020afge	SN IIn/SLSN?	Faint host. FLEET = 50% SN II. Tagged by AD model.
ZTF20acyroio	2021ahwp	SLSN?	Host-$z = 0.545$. Peak $M_{\rm abs} \sim -22.8$ mag.
ZTF21aaahiba	2021ahwq	SN Ia?	⋯
ZTF21aaaolli	2021ahwr	SN Ia?	⋯
ZTF21aaaxnnv	2021ahws	SN I?	⋯
ZTF21aaayfre	2021ahwt	SN Ia?	⋯
ZTF21aabyoxk	2021ahwu	TDE/SN IIn?	Blue. Unclear host. Possible host-$z = 0.14$.
ZTF21aacuckk	2021ahwv	SN I?	Host-$z = 0.094$. Peak $M_{\rm abs} \sim -19.0$ mag.
ZTF21aadplfw	2021ahww	SN Ia?	⋯
ZTF21aadrrbf	2021ahwx	SN Ia?	⋯
ZTF21aadruuo	2021ahwy	SN II?	Edge-on host galaxy.
ZTF21aadsosc	2021ahwz	SN?	No decline.
ZTF21aaekerb	2021ahxa	SN?	⋯
ZTF21aaewjmq	2021ahxb	SN II?	⋯
ZTF21aagkynz	2021ahxc	SN?	⋯
ZTF21aagtspb	2021ahxd	SN I?	⋯
ZTF21aagywat	2021ahxe	SN Ia?	Faint host.
ZTF21aahhegg	2021ahxf	SN Ia?	⋯
ZTF21aaiahsu	2021ahxg	SN Ia-91bg-like?	Host-$z = 0.088$. SALT3 $c = +0.3$. Peak $M_{\rm abs} \sim -18.8$ mag.
ZTF21aaiaqvc	2021ahxh	SN Ia?	⋯
ZTF21aaiqdbp	2021ahxi	SN Ia?	⋯
ZTF21aaiqdkm	2021ahxj	SN Ia?	⋯
ZTF21aaiqifj	2021ahxk	SN Ia?	No visible host.
ZTF21aaitqpu	2021ahxl	SN I?	Faint host.
ZTF21aakitay	2021ahxm	SN Ia?	⋯
ZTF21aakjxhg	2021ahxn	SN Ia?	Ring host galaxy?
ZTF21aalgiex	2021ahxo	SN Ia?	⋯
ZTF21aalimxp	2021ahxp	SN Ib/c?	Tagged by AD model.
ZTF21aantmww	2021ahxq	SN II/IIP?	⋯
ZTF21aaoijsw	2021ahxr	SN Ia?	⋯
ZTF21aapjzyl	2021ahxs	SN Ia?	⋯
ZTF21aaqgmfz	2021ahxt	SN Ia?	⋯
ZTF21aarbdjl	2021ahxu	SN Ia?	⋯
ZTF21aardvtr	2021ahxv	SN Ia?	⋯
ZTF21aarffuz	2021ahxw	SN I?	⋯
ZTF21aarhnil	2021ahxx	SN I?	⋯
ZTF21aarnxig	2021ahxy	SN Ia?	Faint host.
ZTF21aarspet	2021ahxz	SN Ia?	⋯
ZTF21aarvysz	2021ahya	SN?	No visible host.
ZTF21aasjilo	2021ahyb	SN?	⋯
ZTF21aasjlxg	2021ahyc	SN Ia?	⋯
ZTF21aasttsm	2021ahyd	SN Ia?	⋯
ZTF21aatbbjt	2021ahye	SN Ia?	⋯
ZTF21aathago	2021ahyf	SN II/IIP?	Faint host.
ZTF21aatwoxq	2021ahyg	SN I?	⋯
ZTF21aavdtcy	2021ahyh	SN I?	⋯
ZTF21aawlnxl	2021ahyi	SN Ia?	⋯
ZTF21aawzmne	2021ahyj	SN Ia?	Blue. Off-nuclear. SALT3 $x = +3.00$, $c = -0.27$.
ZTF21aaxslhj	2021ahyk	SN?	⋯
ZTF21aaxsqht	2021ahyl	SN Ia?	⋯
ZTF21aaxszmx	2021ahym	SN Ia?	Host-$z = 0.12$. Peak $M_{\rm abs} \sim -19.0$ mag.
ZTF21aayebcv	2021ahyn	SN Ia?	Tagged by AD model.
ZTF21aayebwx	2021ahyo	SN Ia?	Host-$z = 0.148$. Peak $M_{\rm abs} \sim -19.7$ mag.
ZTF21aayfqtb	2021ahyp	SN Ia?	⋯
ZTF21aaynwgp	2021ahyq	SN Ia?	⋯
ZTF21aayotul	2021ahyr	SN Ia?	⋯
ZTF21aazqpgb	2021ahys	SN Ia?	⋯
ZTF21aazzciu	2021ahyt	SN?	⋯
ZTF21abamprv	2021ahyu	SN II?	⋯
ZTF21abbwxbf	2021ahyv	SN Ia?	⋯
ZTF21abcnhrd	2021ahyw	SN Ia?	⋯
ZTF21abcopnn	2021ahyx	SN Ia?	⋯
ZTF21abcothc	2021ahyy	SN Ia?	⋯
ZTF21abcsifi	2021ahyz	SN?	Faint host.
ZTF21abcskis	2021ahza	SN?	⋯
ZTF21abcsqmp	2021ahzb	SN?	⋯
ZTF21abcudam	2021ahzc	SN Ia?	⋯
ZTF21abejeyz	2021ahzd	SN?	No decline.
ZTF21abenlno	2021ahze	SN?	Faint host.

Table F2
(Continued)

ZTF ID	IAU Name	Possible Class	Remarks
ZTF21abezpqc	2021ahzf	SN II/IIP?	⋯
ZTF21abfreru	2021ahzg	SN Ia?	⋯
ZTF21abfspnh	2021ahzh	SN Ia?	⋯
ZTF21abhsttv	2021ahzi	SN Ia?	Faint host.
ZTF21abhzape	2021ahzj	SN Ia?	Host-$z = 0.145$. Peak $M_{\rm abs} \sim -19.4$ mag.
ZTF21abidgad	2021ahzk	SN Ia?	⋯
ZTF21abigfjv	2021ahzl	SN?	Faint host. Tagged by AD model.
ZTF21abjkdli	2021ahzm	SN Ia?	⋯
ZTF21abjqcmb	2021ahzn	SN Ia?	Host-$z = 0.166$. Peak $M_{\rm abs} \sim -19.6$ mag.
ZTF21abjtxur	2021ahzo	SN?	No decline.
ZTF21abkapuy	2021ahzp	SN II?	⋯
ZTF21abowuoq	2021ahzq	SN Ia?	Faint host.
ZTF21abtoqij	2021ahzr	SN I?	⋯
ZTF21abutvjt	2021ahzs	SN Ia?	⋯
ZTF21abvqzuu	2021ahzt	SN II?	⋯
ZTF21abvtzbv	2021ahzu	SN Ia?	⋯
ZTF21abvubre	2021ahzv	SN II?	Faint host.
ZTF21abvufbg	2021ahzw	SN Ia?	Faint host.
ZTF21abwulxy	2021ahzx	SN II?	⋯
ZTF21abwymva	2021ahzy	SN?	Faint host.
ZTF21abxabce	2021ahzz	SN Ia?	Faint host.
ZTF21abxayqz	2021aiaa	SN II?	⋯
ZTF21abxcqfx	2021aiab	SN Ia?	Faint host.
ZTF21abxjdtn	2021aiac	SN Ia?	Faint host.
ZTF21abxkqry	2021aiad	SN Ia?	No visible host.
ZTF21abxoxhb	2021aiae	SN Ia?	⋯
ZTF21abxyydw	2021aiaf	SN II/IIP?	⋯
ZTF21abxzzel	2021aiag	SN Ia?	Faint host.
ZTF21abyaetw	2021aiah	SN Ia?	Faint host.
ZTF21acblsec	2021aiai	SN Ia?	⋯
ZTF21aceylki	2021aiaj	SN Ia?	⋯
ZTF21acguaop	2021aiak	SN Ia?	Edge-on host. Host-$z = 0.096$. Peak $M_{\rm abs} > -18.5$ mag. No rise.
ZTF21achsdxh	2021aial	SN Ia?	⋯
ZTF21acjotbc	2021aiam	SN Ia?	⋯
ZTF21acjouhg	2021aian	SN I?	⋯

Notes. Bold text indicates an object we consider a likely anomaly.
[a] This object is also listed as ZTF18adjzsqt.

We do not include objects discovered previously through our ANN search (see Table F1).

ORCID iDs

P. D. Aleo https://orcid.org/0000-0002-6298-1663
A. W. Engel https://orcid.org/0000-0003-2348-483X
G. Narayan https://orcid.org/0000-0001-6022-0484
C. R. Angus https://orcid.org/0000-0002-4269-7999
K. Malanchev https://orcid.org/0000-0001-7179-7406
K. Auchettl https://orcid.org/0000-0002-4449-9152
V. F. Baldassare https://orcid.org/0000-0003-4703-7276
A. Berres https://orcid.org/0000-0002-5010-441X
T. J. L. de Boer https://orcid.org/0000-0001-5486-2747
B. M. Boyd https://orcid.org/0000-0002-0622-1117
K. C. Chambers https://orcid.org/0000-0001-6965-7789
K. W. Davis https://orcid.org/0000-0002-5680-4660
N. Esquivel https://orcid.org/0009-0002-2806-9379
D. Farias https://orcid.org/0000-0002-6886-269X
R. J. Foley https://orcid.org/0000-0002-2445-5275
A. Gagliano https://orcid.org/0000-0003-4906-8447
C. Gall https://orcid.org/0000-0002-8526-3963
H. Gao https://orcid.org/0000-0003-1015-5367
S. Gomez https://orcid.org/0000-0001-6395-6702
M. Grayling https://orcid.org/0000-0002-6741-983X
D. O. Jones https://orcid.org/0000-0002-6230-0151
C.-C. Lin https://orcid.org/0000-0002-7272-5129
E. A. Magnier https://orcid.org/0000-0002-7965-2815
K. S. Mandel https://orcid.org/0000-0001-9846-4417
T. Matheson https://orcid.org/0000-0001-6685-0479
S. I. Raimundo https://orcid.org/0000-0002-6248-398X
V. G. Shah https://orcid.org/0009-0009-1590-2318
M. D. Soraisam https://orcid.org/0000-0001-6360-992X
K. M. de Soto https://orcid.org/0000-0002-9886-2834
S. Vicencio https://orcid.org/0009-0009-7129-7538
V. A. Villar https://orcid.org/0000-0002-5814-4061
R. J. Wainscoat https://orcid.org/0000-0002-1341-0952

References

Aleo, P. 2024, TNSCR, 2024-844, 1
Aleo, P., Auchettl, K., & Ransome, C. 2023, TNSCR, 2023-2506, 1
Aleo, P. D., Malanchev, K., Sharief, S., et al. 2023, ApJS, 266, 9
Aleo, P. D., Malanchev, K. L., Pruzhinskaya, M. V., et al. 2022, NewA, 96, 101846
Alves, C. S., Peiris, H. V., Lochner, M., et al. 2022, ApJS, 258, 23
Andreoni, I., Coughlin, M. W., Kool, E. C., et al. 2021, ApJ, 918, 63
Angus, C. 2023, TNSCR, 2023-2466, 1
Angus, C., & Aleo, P. 2023, TNSCR, 2023-2467, 1
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33
Ayala, B., Ramirez, M., Dastidar, R., et al. 2023, TNSCR, 2023-2482, 1
Babu, G. J., & Mahabal, A. 2016, INSR, 84, 506
Baldeschi, A., Miller, A., Stroh, M., Margutti, R., & Coppejans, D. L. 2020, ApJ, 902, 60
Barkhudaryan, L. V., Hakobyan, A. A., Karapetyan, A. G., et al. 2019, MNRAS, 490, 718
Barris, B. J., Tonry, J. L., Blondin, S., et al. 2004, ApJ, 602, 571
Bazin, G., Palanque-Delabrouille, N., Rich, J., et al. 2009, A&A, 499, 653
Beck, R., Szapudi, I., Flewelling, H., et al. 2021, MNRAS, 500, 1633
Bell, E. F., Wolf, C., Meisenheimer, K., et al. 2004, ApJ, 608, 752
Bellinger, C., Drummond, C., & Japkowicz, N. 2016, in Joint European Conf. 9851, on Machine Learning and Knowledge Discovery in Databases (Berlin: Springer), 248
Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, PASP, 131, 018002
Bernhardsson, E., 2018 Annoy: Approximate Nearest Neighbors in C++/Python, https://pypi.org/project/annoy/
Blondin, S., & Tonry, J. L. 2007, ApJ, 666, 1024
Böhm, V., Kim, A. G., & Juneau, S. 2023, MNRAS, 526, 3072
Boone, K. 2019, AJ, 158, 257
Boone, K. 2021, AJ, 162, 275
Boone, K., Aldering, G., Antilogus, P., et al. 2021, ApJ, 912, 71
Breiman, L. 2001, Machine Learning, 45, 5
Brout, D., & Scolnic, D. 2021, ApJ, 909, 26
Burhanudin, U. F., & Maund, J. R. 2023, MNRAS, 521, 1601
Byler, N., Dalcanton, J. J., Conroy, C., & Johnson, B. D. 2017, ApJ, 840, 44
Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, ApJ, 345, 245
Carrasco-Davis, R., Reyes, E., Valenzuela, C., et al. 2021, AJ, 162, 231
Chabrier, G. 2003, PASP, 115, 763
Chambers, K. C., Boer, T. D., Bulger, J., et al. 2023, TNSTR, 2023-7501, 1
Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv:1612.05560
Charnock, T., & Moss, A. 2017, ApJL, 837, L28
Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. 2002, Journal of Artificial Intelligence Research, 16, 321
Chiang, Y.-K. 2023, ApJ, 958, 118

Choi, J., Dotter, A., Conroy, C., et al. 2016, ApJ, 823, 102
Chu, M., Dahiwale, A., & Fremling, C. 2021, TNSCR, 2021-3328, 1
Conroy, C., & Gunn, J. E. 2010, ApJ, 712, 833
Conroy, C., Gunn, J. E., & White, M. 2009, ApJ, 699, 486
Coughlin, M. W., Burdge, K., Duev, D. A., et al. 2021, MNRAS, 505, 2954
Coulter, D. A., Jones, D. O., McGill, P., et al. 2022, YSE-PZ: An Open-source Target and Observation Management System, v0.3.0, Zenodo, doi:10.5281/zenodo.7278430
Coulter, D. A., Jones, D. O., McGill, P., et al. 2023, PASP, 135, 064501
Cranmer, K., Brehmer, J., & Louppe, G. 2020, PNAS, 117, 30055
Csabai, I., Budavári, T., Connolly, A. J., et al. 2003, AJ, 125, 580
Cui, K., Armstrong, D. J., & Feng, F. 2024, ApJS, 274, 29
Dálya, G., Galgóczi, G., Dobos, L., et al. 2018, MNRAS, 479, 2374
Davis, K., Foley, R., Izzo, L., et al. 2023, TNSCR, 2023-2235, 1
Davis, K. W., Taggart, K., Tinyanont, S., et al. 2023, MNRAS, 523, 2530
Demianenko, M., Malanchev, K., Samorodova, E., et al. 2023, A&A, 677, A16
D'Isanto, A., Cavuoti, S., Brescia, M., et al. 2016, MNRAS, 457, 3119
Donoso-Oliva, C., Becker, I., Protopapas, P., et al. 2023, A&A, 670, A54
Dotter, A. 2016, ApJS, 222, 8
Dressler, A., Lynden-Bell, D., Burstein, D., et al. 1987, ApJ, 313, 42
Duev, D. A., & van der Walt, S. J. 2021, arXiv:2111.12142
Etsebeth, V., Lochner, M., Walmsley, M., & Grespan, M. 2024, MNRAS, 529, 732
Fakhouri, H. K., Boone, K., Aldering, G., et al. 2015, ApJ, 815, 58
Foley, R. J., & Mandel, K. 2013, ApJ, 778, 167
Foley, R. J., Papenkova, M. S., Swift, B. J., et al. 2003, PASP, 115, 1220
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306
Förster, F., Cabrera-Vives, G., Castillo-Navarrete, E., et al. 2021, AJ, 161, 242
Förster, F., Muñoz Arancibia, A. M., Reyes-Jainaga, I., et al. 2022, AJ, 164, 195
Fossey, S. J., Cooke, B., Pollack, G., Wilde, M., & Wright, T. 2014, CBET, 3792, 1
Frederick, S., Gezari, S., Graham, M. J., et al. 2021, ApJ, 920, 56
Fremling, C. 2023a, TNSTR, 2023-1858, 1
Fremling, C. 2023b, TNSTR, 2023-1800, 1
Fremling, C. 2023c, TNSTR, 2023-1885, 1
Fremling, C. 2023d, TNSTR, 2023-2306, 1
Fremling, C., Miller, A. A., Sharma, Y., et al. 2020, ApJ, 895, 32
Gagliano, A., Contardo, G., Foreman-Mackey, D., Malz, A., & Aleo, P. 2023, AAS Meeting, 241, 103.02
Gagliano, A., Izzo, L., Kilpatrick, C. D., et al. 2022, ApJ, 924, 55
Gagliano, A., Narayan, G., Engel, A., Carrasco Kind, M. & LSST Dark Energy Science Collaboration 2021, ApJ, 908, 170
Gagliano, A., & Villar, V. A. 2023, arXiv:2312.16687
Galama, T. J., Vreeswijk, P. M., van Paradijs, J., et al. 1998, Natur, 395, 670
Gallazzi, A., Charlot, S., Brinchmann, J., White, S. D. M., & Tremonti, C. A. 2005, MNRAS, 362, 41
Gal-Yam, A. 2019, ARA&A, 57, 305
Gezari, S., Martin, D. C., Forster, K., et al. 2013, ApJ, 766, 60
Giles, D., & Walkowicz, L. 2019, MNRAS, 484, 834
Giles, D. K., & Walkowicz, L. 2020, MNRAS, 499, 524
Gomez, S., Berger, E., Blanchard, P. K., et al. 2020, ApJ, 904, 74
Gomez, S., Berger, E., Blanchard, P. K., et al. 2023a, ApJ, 949, 114
Gomez, S., Hiramatsu, D., Berger, E., Blanchard, P., & Gezari, S. 2023b, TNSCR, 2023-2347, 1
Gomez, S., Hiramatsu, D., Kumar, H., Berger, E., & Blanchard, P. 2023c, TNSCR, 2023-2877, 1
Gomez, S., Villar, V. A., Berger, E., et al. 2023d, ApJ, 949, 113
Graur, O., Bianco, F. B., Modjaz, M., et al. 2017, ApJ, 837, 121
Grayling, M., Thorp, S., Mandel, K. S., et al. 2024, MNRAS, 531, 953
Gupta, R., Muthukrishna, D., & Lochner, M. 2024, arXiv:2403.14742
Gutiérrez, C. P., Anderson, J. P., Hamuy, M., et al. 2017, ApJ, 850, 89
Hachinger, S., Mazzali, P. A., Taubenberger, S., Pakmor, R., & Hillebrandt, W. 2009, MNRAS, 399, 1238
Hahn, C., & Melchior, P. 2022, ApJ, 938, 11
Hambleton, K. M., Bianco, F. B., Street, R., et al. 2023, PASP, 135, 105002
Hansson, K. S. A., Lisker, T., & Grebel, E. K. 2012, MNRAS, 427, 2376
He, H., Bai, Y., Garcia, E. A., & Li, S. 2008, in IEEE Int. Joint Conf. on Neural Networks (IEEE World Congress on Computational Intelligence) (Piscataway, NJ: IEEE), 1322
Henne, V., Pruzhinskaya, M. V., Rosnet, P., et al. 2017, NewA, 51, 43
Heringer, E., van Kerkwijk, M. H., Sim, S. A., & Kerzendorf, W. E. 2017, ApJ, 846, 15
Hložek, R., Ponder, K. A., Malz, A. I., et al. 2023, ApJS, 267, 25
Horne, K. 1986, PASP, 98, 609
Howell, D. A., Sullivan, M., Perrett, K., et al. 2005, ApJ, 634, 1190
Hunter, J. D. 2007, CSE, 9, 90
Irani, I., Prentice, S. J., Schulze, S., et al. 2022, ApJ, 927, 10
Ishida, E. E. O., Kornilov, M. V., Malanchev, K. L., et al. 2021, A&A, 650, A195
Itagaki, K. 2023, TNSTR, 2023-1158, 1
Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, ApJ, 873, 111
Jacobson-Galán, W., Venkatraman, P., Margutti, R., et al. 2022, ApJ, 932, 58
Jacobson-Galán, W. V., Dessart, L., Davis, K. W., et al. 2024, ApJ, 970, 189
Jacobson-Galán, W. V., Dessart, L., Jones, D. O., et al. 2022, ApJ, 924, 15
Jarrett, T. H., Cohen, M., Masci, F., et al. 2011, ApJ, 735, 112
Johansson, J., Thomas, D., Pforr, J., et al. 2013, MNRAS, 435, 1680
Johnson, B. D., Leja, J., Conroy, C., & Speagle, J. S. 2021, ApJS, 254, 22
Jolliffe, I. 2002, Principal Component Analysis (New York: Springer)
Jones, D. O., Foley, R. J., Narayan, G., et al. 2021, ApJ, 908, 143
Kelly, P. L., Hicken, M., Burke, D. L., Mandel, K. S., & Kirshner, R. P. 2010, ApJ, 715, 743
Kelsey, L., Sullivan, M., Wiseman, P., et al. 2023, MNRAS, 519, 3046
Kennamer, N., Ishida, E. E. O., Gonzalez-Gaitan, S., et al. 2020, arXiv:2010.05941
Kenworthy, W. D., Jones, D. O., Dai, M., et al. 2021, ApJ, 923, 265
Kessler, R., Bernstein, J. P., Cinabro, D., et al. 2009, PASP, 121, 1028
Khullar, G., Nord, B., Ćiprijanović, A., Poh, J., & Xu, F. 2022, MLS&T, 3, 04LT04
Kim, D.-W., Protopapas, P., Bailer-Jones, C. A. L., et al. 2014, A&A, 566, A43
Kron, R. G. 1980, ApJS, 43, 305
Kuncarayakti, H., Sollerman, J., Izzo, L., et al. 2023, A&A, 678, A209
Kunkel, W., Madore, B., Shelton, I., et al. 1987, IAUC, 4316, 1
Kwon, K. J., Hahn, C., & Alsing, J. 2023, ApJS, 265, 23
Lanusse, F., Parker, L., Golkar, S., et al. 2024, MNRAS, 531, 4990
Lavrukhina, A., & Malanchev, K. 2023, arXiv:2302.10837
Lawrence, A., Warren, S. J., Almaini, O., et al. 2007, MNRAS, 379, 1599
Leja, J., Carnall, A. C., Johnson, B. D., Conroy, C., & Speagle, J. S. 2019a, ApJ, 876, 3
Leja, J., Johnson, B. D., Conroy, C., van Dokkum, P. G., & Byler, N. 2017, ApJ, 837, 170
Leja, J., Johnson, B. D., Conroy, C., et al. 2019b, ApJ, 877, 140
Leja, J., van Dokkum, P. G., Momcheva, I., et al. 2013, ApJL, 778, L24
Leoni, M., Ishida, E. E. O., Peloton, J., & Möller, A. 2022, A&A, 663, A13
Li, W., Leaman, J., Chornock, R., et al. 2011, MNRAS, 412, 1441
Li, X., Ragosta, F., Clarkson, W. I., & Bianco, F. B. 2022, ApJS, 258, 2
Liang, Y., Melchior, P., Lu, S., Goulding, A., & Ward, C. 2023, AJ, 166, 75
Lidman, C., Martin, B., Brout, D., et al. 2023, TNSCR, 2023-1786, 1
Liu, F. T., Ting, K. M., & Zhou, Z.-H. 2012, ACM Trans. Knowl. Discov. Data, 6, 1
Liu, Y., & Modjaz, M. 2014, arXiv:1405.1437
Liu, Y.-Q., Modjaz, M., Bianco, F. B., & Graur, O. 2016, ApJ, 827, 90
Lochner, M., & Bassett, B. A. 2021, A&C, 36, 100481
Lochner, M., McEwen, J. D., Peiris, H. V., Lahav, O., & Winter, M. K. 2016, ApJS, 225, 31
Lokken, M., Gagliano, A., Narayan, G., et al. 2023, MNRAS, 520, 2887
Lyman, J. D., Taddia, F., Stritzinger, M. D., et al. 2018, MNRAS, 473, 1359
Magnier, E. A., Sweeney, W. E., Chambers, K. C., et al. 2020, ApJS, 251, 6
Mahabal, A., Sheth, K., Gieseke, F., et al. 2017, arXiv:1709.06257
Malanchev, K., 2021 Light-curve: Light Curve Analysis Toolbox, Astrophysics Source Code Library, ascl:2107.001
Malanchev, K. L., Pruzhinskaya, M. V., Korolev, V. S., et al. 2021, MNRAS, 502, 5147
Margutti, R., Metzger, B. D., Chornock, R., et al. 2019, ApJ, 872, 18
Martínez-Galarza, J. R., Bianco, F. B., Crake, D., et al. 2021, MNRAS, 508, 5734
Matheson, T., Stubens, C., Wolf, N., et al. 2021, AJ, 161, 107
McInnes, L., Healy, J., Saul, N., & Großberger, L. 2018, JOSS, 3, 861
McIntosh, M., Bieryla, A., Newton, E. R., et al. 2014, AAS Meeting, 224, 121.17
Modjaz, M., Blondin, S., Kirshner, R. P., et al. 2014, AJ, 147, 99
Modjaz, M., Liu, Y. Q., Bianco, F. B., & Graur, O. 2016, ApJ, 832, 108
Möller, A., & de Boissière, T. 2020, MNRAS, 491, 4277
Möller, A., Peloton, J., Ishida, E. E. O., et al. 2021, MNRAS, 501, 3272
Moré, J. J. 1978, The Levenberg-Marquardt Algorithm: Implementation and Theory, Vol. 630 (Berlin: Springer), 105
Moreno-Cartagena, D., Cabrera-Vives, G., Protopapas, P., et al. 2023, arXiv:2308.06404
Munoz-Arancibia, A., Hernandez-Garcia, L., Bauer, F. E., et al. 2023a, TNSAN, 138, 1

Munoz-Arancibia, A., Pignata, G., Bauer, F. E., et al. 2023b, TNSTR, 2023-2237, 1

Munoz-Arancibia, A., Pignata, G., Forster, F., et al. 2023c, TNSTR, 2023-1714, 1

Muthukrishna, D., Mandel, K. S., Lochner, M., Webb, S., & Narayan, G. 2022, MNRAS, 517, 393

Muthukrishna, D., Narayan, G., Mandel, K. S., Biswas, R., & Hložek, R. 2019, PASP, 131, 118002

Nair, P. B., & Abraham, R. G. 2010, ApJS, 186, 427

Narayan, G. & ELAsTiCC Team 2023, AAS Meeting, 241, 117.01

NOIRLab CSDC Team, 2023 ANTARES Client, GitLab, https://gitlab.com/nsf-noirlab/csdc/antares/client/

Nordin, J., Brinnel, V., van Santen, J., et al. 2019, A&A, 631, A147

Nugent, A. E., Fong, W., Dong, Y., et al. 2020, ApJ, 904, 52

Nugent, A. E., Polin, A. E., & Nugent, P. E. 2023, arXiv:2304.10601

Nun, I., Pichara, K., Protopapas, P., & Kim, D.-W. 2014, ApJ, 793, 23

Nun, I., Protopapas, P., Sim, B., & Chen, W. 2016, AJ, 152, 71

Nyholm, A., Sollerman, J., Tartaglia, L., et al. 2020, A&A, 637, A73

Ochsenbein, F., Bauer, P., & Marcout, J. 2000, A&AS, 143, 23

Papamakarios, G., & Murray, I. 2016, arXiv:1605.06376

Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., & Lakshminarayanan, B. 2021, JMLR, 22, 1

Pasquet, J., Pasquet, J., Chaumont, M., & Fouchez, D. 2019, A&A, 627, A21

Patterson, M. T., Bellm, E. C., Rusholme, B., et al. 2019, PASP, 131, 018001

Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, JMLR, 12, 2825, http://jmlr.org/papers/v12/pedregosa11a.html

Perez-Carrasco, M., Cabrera-Vives, G., Hernandez-García, L., et al. 2023, AJ, 166, 151

Perley, D. A., Sollerman, J., Schulze, S., et al. 2022, ApJ, 927, 180

Perley, D. A., Yan, L., Lunnan, R., et al. 2021, TNSAN, 198, 1

Perlmutter, S., Gabi, S., Goldhaber, G., et al. 1997, ApJ, 483, 565

Pierel, J. D. R., Arendse, N., Ertl, S., et al. 2023, ApJ, 948, 115

Pierel, J. D. R., Rodney, S., Avelino, A., et al. 2018, PASP, 130, 114504

Pimentel, Ó., Estévez, P. A., & Förster, F. 2023, AJ, 165, 18

Popovic, B., Brout, D., Kessler, R., Scolnic, D., & Lu, L. 2021, ApJ, 913, 49

Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al. 2018, AJ, 156, 123

Pruzhinskaya, M. V., Malanchev, K. L., Kornilov, M. V., et al. 2019, MNRAS, 489, 3591

Qin, Y.-J., Zabludoff, A., Kisley, M., et al. 2022, ApJS, 259, 13

Qu, H., Sako, M., Möller, A., & Doux, C. 2021, AJ, 162, 67

Quimby, R. M., De Cia, A., Gal-Yam, A., et al. 2018, ApJ, 855, 2

Rasmussen, C. E., & Williams, C. K. I. 2005, Gaussian Processes for Machine Learning (Cambridge, MA: MIT Press)

Rebbapragada, U., Protopapas, P., Brodley, C. E., & Alcock, C. 2009, in ASP Conf. Ser. 411, Astronomical Data Analysis Software and Systems XVIII, ed. D. A. Bohlender, D. Durand, & P. Dowler (San Francisco, CA: ASP), 264

Reyes, E., & Estévez, P. A. 2020, arXiv:2005.07779

Reyes-Jainaga, I., Förster, F., Muñoz Arancibia, A. M., et al. 2023, ApJL, 952, L43

Richardson, D., Jenkins, Robert, L. I., Wright, J., & Maddox, L. 2014, AJ, 147, 118

Russeil, E., Malanchev, K. L., Aleo, P. D., et al. 2024, A&A, 683, A251

Sánchez, P., Lira, P., Cartier, R., et al. 2017, ApJ, 849, 110

Sánchez-Sáez, P., Reyes, I., Valenzuela, C., et al. 2021, AJ, 161, 141

Schmidt, B. P., Suntzeff, N. B., Phillips, M. M., et al. 1998, ApJ, 507, 46

Schulze, S., Yaron, O., Sollerman, J., et al. 2021, ApJS, 255, 29

Sharma, Y., Sollerman, J., Fremling, C., et al. 2023, ApJ, 948, 52

Shi, T., & Horvath, S. 2006, JCGS, 15, 118

Siebert, M. R., Dimitriadis, G., Polin, A., & Foley, R. J. 2020, ApJL, 900, L27

Silverman, J. M., Foley, R. J., Filippenko, A. V., et al. 2012, MNRAS, 425, 1789

Sit, T., Kasliwal, M., Tzanidakis, A., De, K. & ZTF Collaboration 2022, AAS Meeting, 240, 131.03

Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, AJ, 131, 1163

Solarz, A., Bilicki, M., Gromadzki, M., et al. 2017, A&A, 606, A39

Soraisam, M., Matheson, T., Lee, C.-H., et al. 2022, ApJL, 926, L11

Soraisam, M. D., Saha, A., Matheson, T., et al. 2020, ApJ, 892, 112

Stein, G., Harrington, P., Blaum, J., Medan, T., & Lukic, Z. 2021, arXiv:2110.13151

Stern, D., Assef, R. J., Benford, D. J., et al. 2012, ApJ, 753, 30

Stetson, P. B. 1996, PASP, 108, 851

Storey-Fisher, K., Huertas-Company, M., Ramachandra, N., et al. 2021, MNRAS, 508, 2946

Strateva, I., Ivezić, Ž., Knapp, G. R., et al. 2001, AJ, 122, 1861

Sullivan, M., Conley, A., Howell, D. A., et al. 2010, MNRAS, 406, 782

Tachibana, Y., & Miller, A. A. 2018, PASP, 130, 128001

The pandas development team 2020, pandas-dev/pandas: Pandas, v1.2.4, Zenodo, doi:10.5281/zenodo.3509134

Tonry, J., Denneau, L., Weiland, H., et al. 2023, TNSTR, 2023-1398, 1

Tonry, J. L., Stubbs, C. W., Lykke, K. R., et al. 2012, ApJ, 750, 99

Trakhtenbrot, B., Arcavi, I., Ricci, C., et al. 2019, NatAs, 3, 242

Tucker, M. A. 2021, TNSCR, 2021-2720, 1

van der Maaten, L., & Hinton, G. 2008, JMLR, 9, 2579

van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, CSE, 13, 22

van Roestel, J., Duev, D. A., Mahabal, A. A., et al. 2021, AJ, 161, 267

van Velzen, S., Gezari, S., Cenko, S. B., et al. 2019, ApJ, 872, 198

Villar, V. A., Berger, E., Miller, G., et al. 2019, ApJ, 884, 83

Villar, V. A., Cranmer, M., Berger, E., et al. 2021, ApJS, 255, 24

Villar, V. A., Hosseinzadeh, G., Berger, E., et al. 2020, ApJ, 905, 94

Vincenzi, M., Sullivan, M., Firth, R. E., et al. 2019, MNRAS, 489, 5802

Wang, B., Leja, J., Villar, V. A., & Speagle, J. S. 2023, ApJL, 952, L10

Webb, S., Lochner, M., Muthukrishna, D., et al. 2020, MNRAS, 498, 3077

Williamson, M., Modjaz, M., & Bianco, F. B. 2019, ApJL, 880, L22

Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, AJ, 140, 1868

Yao, Y. 2023, TNSCR, 2023-2004, 1